

Too Much Data: Prices and Inefficiencies in Data Markets*

Daron Acemoglu[†] Ali Makhdoumi[‡] Azarakhsh Malekian[§] Asu Ozdaglar[¶]

Abstract

When a user shares her data with an online platform, she typically reveals relevant information about other users. We model a data market in the presence of this type of externality in a setup where one or multiple platforms estimate a user's type with data they acquire from all users and (some) users value their privacy. We demonstrate that the data externalities depress the price of data because once a user's information is leaked by others, she has less reason to protect her data and privacy. These depressed prices lead to excessive data sharing. We characterize conditions under which shutting down data markets improves (utilitarian) welfare. Competition between platforms does not redress the problem of excessively low price for data and too much data sharing, and may further reduce welfare. We propose a scheme based on mediated data-sharing that improves efficiency.

Keywords: data, informational externalities, online markets, privacy.

JEL Classification: D62, L86, D83.

*We are grateful to Alessandro Bonatti and Hal Varian for useful conversations and comments. We gratefully acknowledge financial support from Google, Microsoft, the National Science Foundation, and the Toulouse Network on Information Technology.

[†]Department of Economics, Massachusetts Institute of Technology daron@mit.edu

[‡]Fuqua School of Business, Duke University ali.makhdoumi@duke.edu

[§]Rotman School of Management, University of Toronto azarakhsh.malekian@rotman.utoronto.ca

[¶]Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology asuman@mit.edu

1 Introduction

The data of billions of individuals are currently being utilized for personalized advertising or other online services.¹ The use and transaction of individual data are set to grow exponentially in the coming years with more extensive data collection from new online apps and integrated technologies such as the Internet of Things and with the more widespread applications of artificial intelligence (AI) and machine learning techniques. Most economic analyses emphasize benefits from the use and sharing of data because this permits better customization, better information, and more input into AI applications. It is often claimed that because data enables a better allocation of resources and more or higher quality innovation, the market mechanism generates too little data sharing (e.g., [Varian \[2009\]](#), [Jones et al. \[2018\]](#), [Farboodi et al. \[2019\]](#), and [Veldkamp and Chung \[2019\]](#)). Economists have recognized that consumers might have privacy concerns (e.g., [Stigler \[1980\]](#), [Posner \[1981\]](#), and [Varian \[2009\]](#)), but have often argued that data markets could appropriately balance privacy concerns and the social benefits of data (e.g., [Laudon \[1996\]](#) and [Posner and Weyl \[2018\]](#)). In any case, the willingness of the majority of users to allow their data to be used for no or very little direct benefits is argued to be evidence that most users place only a small value on privacy.²

This paper, in contrast, argues that there are forces that will make individual-level data underpriced and the market economy generate too much data. The reason is simple: when an individual shares her data, she compromises not only her own privacy but also the privacy of other individuals whose information is correlated with hers. This negative externality tends to create excessive data sharing. Moreover, when there is excessive data sharing, each individual will overlook her privacy concerns and part with her own information because others' sharing decisions will have already revealed much about her.

The following example illustrates the nature of the problem, introduces some of our key concepts, and clarifies why there will be excessive data sharing and very little willingness to protect privacy on the part of users. Consider a platform with two users, $i = 1, 2$. Each user owns her own personal data, which we represent with a random variable X_i (from the viewpoint of the platform). The relevant data of the two users are related, which we capture by assuming that their random variables are jointly normally distributed with mean zero and correlation coefficient ρ . The platform can acquire or buy the data of a user in order to better estimate her preferences or actions. Its objective is to minimize the mean square error of its estimates of user types, or maximize the amount of *leaked information* about them. Suppose that the valuation (in monetary terms) of the platform for the users' leaked information is one, while the value that the first user attaches to her privacy, again in terms of leaked information about her, is $1/2$ and for the second user it is $v > 0$. We also assume that the platform makes take-it-or-leave-it offers to the users to purchase

¹Just Facebook has almost 2.5 billion monthly (active) individual users.

²Consumers often report valuing privacy (e.g., [Westin \[1968\]](#); [Goldfarb and Tucker \[2012\]](#)), but do not take much action to protect their privacy (e.g., ["Why your inbox is crammed full of privacy policies"](#), [WIRED](#), May 24, 2018 and [Athey et al. \[2017\]](#)).

their data. In the absence of any restrictions on data markets or transaction costs, the first user will always sell her data (because her valuation of privacy, $1/2$, is less than the value of information to the platform, 1). But given the correlation between the types of the two users, this implies that the platform will already have a fairly good estimate of the second user's information. Suppose, for illustration, that $\rho \approx 1$. In this case, the platform will know almost everything relevant about user 2 from user 1's data, and this undermines the willingness of user 2 to protect her data. In fact, since user 1 is revealing almost everything about her, she would be willing to sell her own data for a very low price (approximately 0 given $\rho \approx 1$). But once the second user is selling her data, this also reveals the first user's data, so the first user can only charge a very low price for her data. Therefore in this simple example, the platform will be able to acquire both users' data at approximately zero price. Critically, however, this price does not reflect the users' valuation of privacy. When $v \leq 1$, the equilibrium is efficient because data are socially beneficial in this case (even if data externalities change the distribution of economic surplus to the advantage of the platform). However, it can be arbitrarily inefficient when v is sufficiently high. This is because the first user, by selling her data, is creating a negative externality on the second user.

We develop a stylized and tractable reduced-form model, consisting of a community of users with correlated information, to explore more systematically the ideas illustrated by this example. We analyze the model both under a monopoly platform and under competition between platforms trying to simultaneously attract users and acquire their data.

Our main results correspond to generalizations of the insights summarized by the preceding example. First, we introduce our general framework and characterize the first-best allocation which maximizes the sum of surplus of users and platforms. The first best typically involves considerable data transactions, but those individuals creating significant (negative) externalities on others should not share their data. Second, we establish the existence of an equilibrium and characterize the prices at which data will be transacted. This characterization clarifies how the market price of data for a user and the distribution of surplus depend on information leaked by other users. Third and more importantly, we provide conditions under which the equilibrium in the data market is inefficient as well as conditions for simple restrictions on markets to improve welfare. At the root of these inefficiencies are the economic forces already highlighted by our example: inefficiencies arise when a subset of users are willing to part with their data, which are informative about other users whose value of privacy is high. We show that these insights extend to environments with competing platforms and incomplete information as well.

We further investigate various policy approaches to data markets. Person-specific taxes on data transactions can restore the first best, but are impractical. We show in addition how uniform taxation on all data transactions might, under some conditions, improve welfare. Finally, we propose a new regulation scheme where data transactions are mediated in a way that reduces their correlation with the data of other users, thus minimizing leaked information about others. We additionally develop a procedure for implementing this scheme based on "*de-correlation*", meaning

transforming users' data so that their correlation with others' data and types is removed.³

Our paper relates to the literature on privacy and its legal and economic aspects. The classic definition of privacy, proposed by justices Warren and Brandeis in 1890, is the protection of someone's personal space and the right to be let alone (Warren and Brandeis [1890]). Relatedly, and more relevant to our focus, Westin [1968] defines it as the control over and safeguarding of personal information, and this perspective has been explored from various angles in recent work (e.g., Pasquale [2015], Tirole [2019], Zuboff [2019]).

More closely related to our paper include MacCarthy, Boyd [2011], and Fairfield and Engel [2015] who are the first contributions we are aware of that emphasize externalities in data sharing. More recently, Choi et al. [2019] develop a model with a related informational externality and a number of results similar to our excessive information sharing finding. There are several important differences between this paper and ours, however. First, Choi et al. [2019] assume that consumers are identical, while our above example illustrates that heterogeneity in privacy concerns plays a critical role in the inefficiencies in data markets. Our analysis highlights that there are only limited inefficiencies when users are homogeneous (specifically, the equilibrium is efficient in this case when they have low or sufficiently high value of privacy). Second, in contrast to this paper, much of our analysis is devoted to the study of how the correlation structure across different users jointly determines sharing decisions, prices, and the amount of leaked information. Third, their paper does not analyze the case with competing platforms. More recent and independent work by Bergemann et al. [2019] also studies an environment with data externalities. Though there are some parallels between the two papers, their work is different from and largely complementary to ours. In particular, they analyze an economy with symmetric users where there is a monopolist platform and data are used by this monopolist or other downstream firms (such as advertisers) for price discrimination. They also consider learning (willingness to pay) on the users' side and focus on the implications for market prices, profits, and efficiency of the structure of the downstream market and whether data are collected in an anonymized or non-anonymized form. Other recent and relevant contributions to this literature include Fainmesser et al. [2019] and Jullien et al. [2020], which consider the negative effects of leaking user's (private) personalized data but do not study data externalities; Gradwohl [2017] which investigates user behavior in the presence of data externality but does not analyze prices and inefficiencies; and Ichihashi [2019], Ichihashi [2020b] and Ichihashi [2020a] which study the role of information intermediaries and dynamic data collection by platforms.

Our paper also relates to the growing literature on information markets. One branch of this literature focuses on the use of personal data for improved allocation of online resources (e.g., Bergemann and Bonatti [2015], Goldfarb and Tucker [2011], and Montes et al. [2019]). Another branch investigates how information can be monetized either by dynamic sales or optimal mechanisms. For example, Anton and Yao [2002], Babaioff et al. [2012], Esó and Szentes [2007], Horner

³This de-correlation procedure is different from anonymization of data because it does not hide information about the user sharing her data but about others who are correlated with this user.

and Skrzypacz [2016], Bergemann et al. [2018], and Eliaz et al. [2019] consider either static or dynamic mechanisms for selling data, Ghosh and Roth [2015] uses differential privacy framework of Dwork et al. [2014] and study mechanism design with privacy constraints, and Admati and Pfleiderer [1986] and Begenau et al. [2018] study markets for financial data. A third branch focuses on optimal collection and acquisition of information, for example, Agarwal et al. [2019], Chen and Zheng [2019], and Chen et al. [2018]. Lastly, a number of papers investigate the question of whether information harms consumers, either because users are unaware of the data being collected about them (Taylor [2004]) or because of price discrimination related reasons (Acquisti and Varian [2005]). See Acquisti et al. [2016], Bergemann and Bonatti [2019], and Agrawal et al. [2018] for excellent surveys of different aspects of this literature.

The rest of the paper proceeds as follows. Section 2 presents our model, focusing on the case with a single platform for simplicity. Section 3 provides our main results, in particular, characterizing the structure of equilibria in data markets and highlighting their inefficiency due to data externalities. It also shows how shutting down data markets may improve welfare. Section 4 extends these results to a setting with competing platforms, while Section 5 presents a number of generalizations. Section 6 studies how taxes and third-party-mediated information sharing schemes can improve welfare. Section 7 concludes, while Appendix A presents the proofs of some of the results stated in the text and the online Appendix contains the remaining proofs and additional results.

2 Model

In this section we introduce our model, focusing first on the case with a single platform. Competition between platforms is analyzed in Section 4.

2.1 Information and Payoffs

We consider n users represented by the set $\mathcal{V} = \{1, \dots, n\}$. Each user $i \in \mathcal{V}$ has a type denoted by x_i which is a realization of a random variable X_i . We assume that the vector of random variables $\mathbf{X} = (X_1, \dots, X_n)$ has a joint normal distribution $\mathcal{N}(0, \Sigma)$, where $\Sigma \in \mathbb{R}^{n \times n}$ is the covariance matrix of \mathbf{X} . Let Σ_{ij} designate the (i, j) -th entry of Σ and $\Sigma_{ii} = \sigma_i^2 > 0$ denote the variance of individual i 's type.

Each user has some personal data, S_i , which are informative about her type. These include both data that user activity on the platform generates (such as from search and purchase histories) and additional data that users may share about their preferences, contacts or past behavior. We suppose that $S_i = X_i + Z_i$ where Z_i is an independent random variable with standard normal distribution, i.e., $Z_i \sim \mathcal{N}(0, 1)$.⁴

⁴For transparency, we assume that both user type and personal data are represented by one-dimensional variables, but all of our main results and insights generalize to a setting with multi-dimensional types and data.

For any user joining the platform, the platform can derive additional revenue if it can predict her type. This might be because of improved personalized services, targeted advertising, or price discrimination for some services sold on the platform. Since the exact source of revenue for the platform is immaterial for our analysis, we simply assume that the platform’s revenue from each user is a(n inverse) function of the mean square error of its forecast of the user’s type, minus what the platform pays to users to acquire their information. Namely, the objective of the platform is to minimize

$$\sum_{i \in \mathcal{V}} \left(\mathbb{E} \left[(\hat{x}_i(\mathbf{S}) - X_i)^2 \right] - \sigma_i^2 + p_i \right), \quad (1)$$

where \mathbf{S} is the vector of data the platform acquires, $\hat{x}_i(\mathbf{S})$ is the platform’s estimate of the user’s type given this information, $-\sigma_i^2$ is included as a convenient normalization, and p_i denotes payments to user i from the platform. This price represents both direct payments to the users in exchange for the type and amount of data shared and indirect payments, for example, in the form of some good or service the platform provides to the user in exchange for her data.

Users value their privacy, which we also model in a reduced-form manner as a function of the same mean square error.⁵ This reflects both pecuniary and nonpecuniary motives, for example, the fact that a user may receive a greater consumer surplus when the platform knows less about her or she may have a genuine demand for keeping her preferences, behavior, and information private. There may also be political and social reasons for privacy, for example, for concealing dissident activities or behaviors disapproved by some groups. We assume, specifically, that user i ’s value of privacy is $v_i \geq 0$, and her payoff is

$$v_i \left(\mathbb{E} \left[(\hat{x}_i(\mathbf{S}) - X_i)^2 \right] - \sigma_i^2 \right) + p_i.$$

This expression and its comparison with (1) clarifies that the platform and users have potentially-opposing preferences over information about user type. We have again subtracted σ_i^2 as a normalization, which ensures that if the platform acquires no additional information about the user and makes no payment to her, her payoff is zero.⁶

Critically, users with $v_i < 1$ value their privacy less than the valuation that the platform attaches to information about them, and thus reducing the mean square error of the estimates of their types is socially beneficial. In contrast, users with $v_i > 1$ value their privacy more, and reducing their mean square error is socially costly. In a world without data externalities (where data about one user have no relevance to the information about other users), the first group of users should allow the platform to acquire (buy) their data, while the second group should not. A simple market mechanism based on prices for data can implement this efficient outcome.

⁵For simplicity, we postpone the introduction of joining decisions to Section 4.

⁶The positive social benefits from data are represented by the platform’s payoff function. This may be because the platform can price its other services in such a way as to capture all of these gains from users. But in our analysis, this assumption is imposed mainly for notational simplicity. If these social benefits from data were shared between the platform and users so that the fraction $\beta_i < 1$ of these gains go directly to users, all of our results would apply without any modification.

We will see that the situation is very different in the presence of data externalities.

2.2 Leaked Information

A key notion for our analysis is *leaked information*, which captures the reduction in the mean square error of the platform's estimate of the type of a user. When the platform has no information about user i , its estimate satisfies $\mathbb{E}[(\hat{x}_i - X_i)^2] = \sigma_i^2$. As the platform receives data from this and other users, its estimate improves and the mean square error declines. The notion of leaked information captures this reduction in mean square error.

Specifically, let $a_i \in \{0, 1\}$ denote the data sharing action of user $i \in \mathcal{V}$ with $a_i = 1$ corresponding to sharing. Denote the profile of sharing decisions by $\mathbf{a} = (a_1, \dots, a_n)$ and the decisions of agents other than i by \mathbf{a}_{-i} . We also use the notation $\mathbf{S}_{\mathbf{a}}$ to denote the data of all individuals for whom $a_j = 1$, i.e., $\mathbf{S}_{\mathbf{a}} = (S_j : j \in \mathcal{V} \text{ s.t. } a_j = 1)$. Given a profile of actions \mathbf{a} , the *leaked information* of (or about) user $i \in \mathcal{V}$ is the reduction in the mean square error of the best estimator of the type of user i :

$$\mathcal{I}_i(\mathbf{a}) = \sigma_i^2 - \min_{\hat{x}_i} \mathbb{E}[(X_i - \hat{x}_i(\mathbf{S}_{\mathbf{a}}))^2].$$

Notably, because of data externalities, leaked information about user i depends not just on her decisions but also on the sharing actions taken by all users. With this notion at hand, we can write the payoff of user i given the price vector $\mathbf{p} = (p_1, \dots, p_n)$ as

$$u_i(a_i, \mathbf{a}_{-i}, \mathbf{p}) = \begin{cases} p_i - v_i \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) & a_i = 1 \\ -v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) & a_i = 0, \end{cases}$$

where recall that $v_i \geq 0$ is user's value of privacy. The platform's objective is to minimize (1) or to maximize

$$U(\mathbf{a}, \mathbf{p}) = \sum_{i \in \mathcal{V}} \mathcal{I}_i(\mathbf{a}) - \sum_{i \in \mathcal{V}: a_i=1} p_i. \quad (2)$$

2.3 Equilibrium Concept

An action profile $\mathbf{a} = (a_1, \dots, a_n)$ and a price vector $\mathbf{p} = (p_1, \dots, p_n)$ constitute a pure strategy equilibrium if both users and the platform maximize their payoffs given other players' strategies. More formally, in the next definition we define an equilibrium as a *Stackelberg equilibrium* in which the platform chooses the price vector recognizing the *user equilibrium* that will result following this choice.

Definition 1. Given the price vector $\mathbf{p} = (p_1, \dots, p_n)$, an action profile $\mathbf{a} = (a_1, \dots, a_n)$ is user equilibrium if for all $i \in \mathcal{V}$,

$$a_i \in \operatorname{argmax}_{a \in \{0,1\}} u_i(a_i = a, \mathbf{a}_{-i}, \mathbf{p}).$$

We denote the set of user equilibria at a given price vector \mathbf{p} by $\mathcal{A}(\mathbf{p})$. A pair $(\mathbf{p}^E, \mathbf{a}^E)$ of price and action vectors is a pure strategy Stackelberg equilibrium if $\mathbf{a}^E \in \mathcal{A}(\mathbf{p}^E)$ and there is no profitable deviation for the platform, i.e.,

$$U(\mathbf{a}^E, \mathbf{p}^E) \geq U(\mathbf{a}, \mathbf{p}), \quad \text{for all } \mathbf{p} \text{ and for all } \mathbf{a} \in \mathcal{A}(\mathbf{p}).$$

In what follows, we refer to a pure strategy Stackelberg equilibrium simply as an equilibrium. The notion of Stackelberg equilibrium in Definition 1 is a refinement of subgame perfect equilibrium and ensures that the platform can choose the best action profile among those that are best responses for users. Without this refinement, there may be additional multiplicity of equilibria.

3 Analysis

In this section, we first study the first-best information sharing decisions that maximize the sum of users and platform payoffs and then proceed to characterizing the equilibrium and its efficiency properties.

3.1 First Best

We define the first best as the data sharing decisions that maximize utilitarian social welfare or social surplus given by the sum of the payoffs of the platform and users. Social surplus from an action profile \mathbf{a} is

$$\text{Social surplus}(\mathbf{a}) = U(\mathbf{a}, \mathbf{p}) + \sum_{i \in \mathcal{V}} u_i(\mathbf{a}, \mathbf{p}) = \sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}).$$

Prices do not appear in this expression because they are transfers from the platform to users.⁷ The first-best action profile, \mathbf{a}^W , maximizes this expression. The next proposition characterizes the first-best action profile.

Proposition 1. *The first best involves $a_i^W = 1$ if*

$$\sum_{j \in \mathcal{V}} (1 - v_j) \frac{(\text{Cov}(X_i, X_j \mid a_i = 0, \mathbf{a}_{-i}^W))^2}{1 + \sigma_j^2 - \mathcal{I}_j(a_i = 0, \mathbf{a}_{-i}^W)} \geq 0, \quad (3)$$

and $a_i^W = 0$ if the left-hand side of (3) is negative.

The proof of this proposition as well as all other proofs, unless otherwise stated, are presented in Appendix A.

⁷In including the platform's payoff in social surplus we are assuming that this payoff is not coming from shifting revenues from some other (perhaps off-line) businesses. If we do not include the payoff of the platform in our welfare measure, our inefficiency results would hold a fortiori.

To understand this result, consider first the case in which there are no data externalities so that the covariance terms in (3) are zero, except $\text{Cov}(X_i, X_i | a_i = 0, \mathbf{a}_{-i}^W) = \sigma_i^2$, so that the left-hand side is simply $\sigma_i^4 / (1 + \sigma_i^2)$ times $1 - v_i$. This yields $a_i^W = 1$ if $v_i \leq 1$. The situation is different in the presence of data externalities, because now the covariance terms are non-zero. In this case, an individual should optimally share her data only if it does not reveal too much about users with $v_j > 1$.

3.2 Equilibrium Preliminaries

The next lemma characterizes two important properties of the leaked information function $\mathcal{I}_i : \{0, 1\}^n \rightarrow \mathbb{R}$.

Lemma 1. 1. Monotonicity: for two action profiles \mathbf{a} and \mathbf{a}' with $\mathbf{a} \geq \mathbf{a}'$,

$$\mathcal{I}_i(\mathbf{a}) \geq \mathcal{I}_i(\mathbf{a}'), \quad \text{for all } i \in \{1, \dots, n\}.$$

2. Submodularity: for two action profiles \mathbf{a} and \mathbf{a}' with $\mathbf{a}'_{-i} \geq \mathbf{a}_{-i}$,

$$\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) \geq \mathcal{I}_i(a_i = 1, \mathbf{a}'_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}'_{-i}).$$

The monotonicity property states that as the set of users who share their information expands, the leaked information about each user (weakly) increases. This is an intuitive consequence of the fact that more information always facilitates the estimation problem of the platform and reduces the mean square error of its estimates. More important for the rest of our analysis is the submodularity property, which implies that the marginal increase in the leaked information from individual i 's sharing decision is decreasing in the information shared by others. This too is intuitive and follows from the fact that when others' actions reveal more information, there is less to be revealed by the sharing decision of any given individual.

Using Lemma 1 we next show that for any price vector $\mathbf{p} \in \mathbb{R}^n$, the set $\mathcal{A}(\mathbf{p})$ is a (non-empty) complete lattice.

Lemma 2. For any \mathbf{p} , the set $\mathcal{A}(\mathbf{p})$ is a complete lattice, and thus has a least and a greatest element.

Lemma 2 implies that the set of user equilibria is always non-empty, but may not be singleton as we illustrate in the next example.

Example 1. Suppose there are two users 1 and 2 with covariance matrix Σ such that $\Sigma_{11} = \Sigma_{22} = 1$ and $\Sigma_{12} = \Sigma_{21} = \rho$ and values $v_1 = v_2 = 1$. The set of user equilibria in this case is depicted in Figure 1. When $p_1, p_2 \in \left[\frac{(2-\rho^2)^2}{2(4-\rho^2)}, \frac{1}{2} \right]$, both action profiles $a_1 = a_2 = 0$ and $a_1 = a_2 = 1$ are user equilibria. This is a consequence of the submodularity of the leaked information function (Lemma 1): when user 1 shares her data, she is also revealing a lot about user 2, and making it less costly for her to share her data. Conversely, when user 1 does not share, this encourages user 2 not to share. Despite this multiplicity of user equilibria, there exists a unique (Stackelberg) equilibrium




Figure 1: The user equilibrium as a function of price vector (p_1, p_2) in the setting of Example 1. For the prices in the purple area in the center, both $a_1 = a_2 = 0$ and $a_1 = a_2 = 1$ are user equilibria.

for this game given by $a_1^E = a_2^E = 1$ and $p_1^E = p_2^E = \frac{(2-\rho^2)^2}{2(4-\rho^2)}$. This uniqueness follows because the platform can choose the price vector to encourage both users to share.

3.3 Existence of Equilibrium

The next theorem establishes the existence of a (pure strategy) equilibrium.

Theorem 1. *An equilibrium always exists. That is, there exist an action profile \mathbf{a}^E and a price vector \mathbf{p}^E such that $\mathbf{a}^E \in \mathcal{A}(\mathbf{p}^E)$, and*

$$U(\mathbf{a}^E, \mathbf{p}^E) \geq U(\mathbf{a}, \mathbf{p}), \quad \text{for all } \mathbf{p} \text{ and for all } \mathbf{a} \in \mathcal{A}(\mathbf{p}). \quad (4)$$

Note that the equilibrium may not be unique, but if there are multiple equilibria, all of them yield the same payoff for the platform (since otherwise (4) would not be satisfied for the equilibrium with lower payoff for the platform).

3.4 An Illustrative Example

In this subsection, we provide an illustrative example that highlights a few of the subtle aspects of the equilibrium. Consider the same setting as in Example 1 with two users with the same value of privacy, v , and a correlation coefficient ρ between their information. We first show that the total payment from the platform to users is non-monotone in the number of users sharing their information. When the platform induces both users to share ($a_1 = a_2 = 1$), it makes a total payment of $v \frac{(2-\rho^2)^2}{4-\rho^2}$. In contrast, when it only induces the first user to share ($a_1 = 1, a_2 = 0$), this will cost $\frac{v}{2}$. Therefore, when $\rho^2 \geq \frac{7-\sqrt{17}}{4} \approx 0.71$, the platform pays less to have both users share their data. Intuitively, this cost-saving for the platform is a consequence of the submodularity of leaked information (Lemma 1): when both users share, the data of each are less valuable in view




Figure 2: Equilibrium and social surplus as a function of the value of privacy v for a setting with two users with $\sigma_1^2 = \sigma_2^2 = 1$, $\Sigma_{12} = \rho$, and $v_1 = v_2 = v$.

of the information revealed by the other user. This finding reflects one of the claims made in the Introduction: market prices for data do not reflect the value that users attached to their privacy and may be depressed because of data externalities.

We next illustrate that equilibrium (social) surplus is non-monotonic in the users' value of privacy. Equilibrium surplus is depicted in Figure 2. For values of v larger than $\frac{4}{(2-\rho^2)^2}$, users do not share their data and equilibrium surplus is zero. When v is smaller than 1, users share their data and equilibrium surplus is positive. For intermediate values of v , in particular for $v \in [1, \frac{4}{(2-\rho^2)^2}]$, the platform chooses a price vector that induces both users to share their data, but in this case, the social surplus is negative. The intuition is related to the point already emphasized in the previous paragraph: when both users share their data, the externalities depress the market prices for data and this makes it profitable for the platform to acquire the users' data even though $v > 1$. More explicitly, when user 2 shares her data, this reveals sufficient information about user 1 that she becomes willing to accept a relatively low price for sharing her data, and this maintains an equilibrium with low prices for data even though both users attach a relatively high value to their privacy.

3.5 Equilibrium Prices

In this subsection, we characterize the equilibrium price vector. For any action profile $\mathbf{a} \in \{0, 1\}^n$, let $\mathbf{p}^{\mathbf{a}}$ denote the least (element-wise minimum) equilibrium price vector that sustains an action profile \mathbf{a} in a user equilibrium. More specifically, $\mathbf{p}^{\mathbf{a}}$ is defined such that:⁸

$$\mathbf{p}^{\mathbf{a}} \leq \mathbf{p}, \quad \text{for all } \mathbf{p} \text{ such that } \mathbf{a} \in \mathcal{A}(\mathbf{p}).$$

⁸Prices for users not sharing their data are not well-defined.

Profit maximization by the platform implies that equilibrium prices must satisfy this property — since otherwise the platform could reduce prices and still implement the same action profile. We therefore refer to $\mathbf{p}^{\mathbf{a}}$ as “equilibrium price vector” or simply as “equilibrium prices” (with the understanding that these would be the equilibrium prices when the platform chooses to induce action profile \mathbf{a}).

The next theorem computes this price vector (and shows that it exists).

Theorem 2. *For any action profile $\mathbf{a} \in \{0, 1\}^n$, we have*

$$\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) = \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) + \frac{(\sigma_i^2 - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}))^2}{(\sigma_i^2 + 1) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})}, \quad (5)$$

and

$$\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) = \mathbf{d}_i^T (I + D_i)^{-1} \mathbf{d}_i, \quad \text{for all } a_i = 1,$$

where D_i is the matrix obtained by removing row and column i from matrix Σ as well as all rows and columns j for which $a_j = 0$, and \mathbf{d}_i is $(\Sigma_{ij} : j \text{ s.t. } a_j = 1)$. The equilibrium price that sustains action profile \mathbf{a} is

$$p_i^{\mathbf{a}} = \begin{cases} v_i \frac{(\sigma_i^2 - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}))^2}{(\sigma_i^2 + 1) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})} & a_i = 1 \\ 0 & a_i = 0. \end{cases}$$

The first part of Theorem 2 provides a decomposition of leaked information about user i in terms of leaked information about her when she does not share her data. In particular, the first term on the right-hand side of the equation (5), $\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$, is her leaked information resulting from the data sharing of other users and thus represents the data externality. The second term is the additional leakage when user i shares her data. The second part of Theorem 2 states that, because the platform offers the prices, the equilibrium price for any user i who shares her information must equal her reservation value, making her indifferent between sharing and not sharing. This result explains why the equilibrium price, $p_i^{\mathbf{a}}$, is equal to the value of privacy, v_i , multiplied by the second term in (5), which is the additional leakage of information and hence the loss of privacy resulting from the user’s own data sharing.

The following is an immediate corollary of Theorem 2.

Corollary 1. *For any user i , the equilibrium price $p_i^{(a_i=1, \mathbf{a}_{-i})}$ (that induces $a_i = 1$ for any action profile $\mathbf{a}_{-i} \in \{0, 1\}^{n-1}$) is increasing in σ_i^2 and decreasing in the data externality captured by $\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$. Moreover, leaked information $\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i})$ is increasing in σ_i^2 and in the data externality $\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$.*

The first part of Corollary 1 shows that a higher variance of user’s type, σ_i^2 , increases the equilibrium price. Intuitively, a higher variance makes the user’s type more difficult to predict and thus her own information more valuable. This also explains why the price is decreasing in the data externality — represented by information leaked by others, $\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$. The last part of Corollary 1 shows that a higher variance of individual type, as well as a greater data externality, increase the overall leakage of information about the user.

The next proposition establishes that equilibrium prices are nonincreasing in the set of users that share their data as well.

Proposition 2. *For two action profiles \mathbf{a}, \mathbf{a}' with $\mathbf{a}' \geq \mathbf{a}$, we have $p_i^{\mathbf{a}'} \leq p_i^{\mathbf{a}}$ for all $i \in \mathcal{V}$ for which $a_i = 1$.*

Proposition 2 follows from Theorem 2 and Lemma 1. In particular, using Theorem 2 the equilibrium price for user i is her additional loss of privacy (increase in the information leakage multiplied by v_i) if she shares her data. From the submodularity of information leakage (Lemma 1) the additional information the user leaks about herself decreases when more people share their data.⁹

3.6 Inefficiency

This subsection presents one of our main results, documenting the extent of inefficiency in data markets.

First note that all users with value of privacy less than 1 will always share their data in equilibrium. For future reference, we state this straightforward result as a lemma.

Lemma 3. *All users with value of privacy $v_i \leq 1$ share their data in equilibrium.*¹⁰

Motivated by this lemma, we partition users into two sets, those with value of privacy below 1 (“low-value users”) and those above (“high-value users”):

$$\mathcal{V}^{(l)} = \{i \in \mathcal{V} : v_i \leq 1\} \text{ and } \mathcal{V}^{(h)} = \{i \in \mathcal{V} : v_i > 1\}.$$

We also denote by $\mathbf{v}^{(h)}$ and $\mathbf{v}^{(l)}$ the vectors of valuations of privacy for high-value and low-value users, respectively. Lemma 3 then implies that for all $i \in \mathcal{V}^{(l)}$ we have $a_i^E = 1$.

The next theorem provides conditions for efficiency and inefficiency. More precisely, we show that if every high-value user is uncorrelated with all other users, then equilibrium is efficient. Otherwise, if either a high-value user is correlated with a low-value user or two high-value users are correlated, there exists a set of valuations (consistent with the set of high and low-value users) such that any equilibrium is inefficient.

Theorem 3. *1. Suppose every high-value user is uncorrelated with all other users. Then the equilibrium is efficient.*

⁹Notice also the connection between Proposition 2 and Crémer and Mclean [1988], who establish in the context of a mechanism design problem with correlated values that when agents reveal more information about each other, then their information rent becomes smaller. In our setting when more users share their data, the value of another user sharing her data becomes small, but this result originates from the correlation in personal data, not from correlated values.

¹⁰The only subtlety here is about users with $v_i = 1$. If these users’ information is correlated with others who are already sharing, their equilibrium price will be strictly less than 1, and this will make it strictly beneficial for the platform to purchase their data. If they are correlated with others who are not sharing, then the platform would still like to purchase these data because of the additional reduction in the mean square error of its estimates of others’ types they enable. When such an individual is uncorrelated with anybody else, then the platform would be indifferent between purchasing her data and not. In this case, for simplicity of notation, we suppose that it still purchases.

2. Suppose at least one high-value user is correlated (has a non-zero correlation coefficient) with a low-value user. Then there exists $\bar{v} \in \mathbb{R}^{|\mathcal{V}^{(h)}|}$ such that for $\mathbf{v}^{(h)} \geq \bar{v}$ the equilibrium is inefficient.
3. Suppose every high-value user is uncorrelated with all low-value users and at least one high-value user is correlated with another high-value user. Let $\tilde{\mathcal{V}}^{(h)} \subseteq \mathcal{V}^{(h)}$ be the subset of high-value users correlated with at least one other high-value user. Then for each $i \in \tilde{\mathcal{V}}^{(h)}$ there exists $\bar{v}_i > 1$ such that if for any $i \in \tilde{\mathcal{V}}^{(h)}$ $v_i < \bar{v}_i$, the equilibrium is inefficient

Theorem 3 clarifies the source of inefficiency in our model. If high-value users are not correlated with others, the equilibrium is efficient. In this case, there may still be data externalities among low-value users and these may affect market prices (and the distribution of economic gains between the users and the platform). But they do not create a loss of privacy for users who prefer not to share their data.

However, the second part of the theorem shows that if high-value users are correlated with low-value users, the equilibrium is typically inefficient. The additional condition $\mathbf{v}^{(h)} \geq \bar{v}$ is not a restrictive one as highlighted in Example 2 below and rules out cases in which high-value users suffer only little loss of privacy but generate socially valuable information about low-value users. In general, the inefficiency identified in this part of the theorem can take one of two forms: either high-value users do not share their data, but because of information leaked about them, they suffer a loss of privacy. Or given the amount of leaked information about them, high-value users decide to share themselves—even though, absent the correlation with low-value users or low-value users' data sharing, they would have preferred not to do so.

Finally, the third part of the theorem covers the remaining case, where high-value users are uncorrelated with low-value users but are correlated among themselves. The equilibrium is again inefficient, because the platform can induce some of them to share their data (even though individually each would prefer not to). This is because when a subset of them share, this compromises the privacy of others, depresses data prices, and may incentivize them to share too (in turn further depressing data prices). This inefficiency applies when some high-value users have intermediate values of privacy (i.e., $v_i \in (1, \bar{v}_i)$), since those with sufficiently high value of privacy cannot be induced to share their data.

Overall, this theorem highlights that inefficiency in data markets originates from the combination of sufficiently high value attached to privacy by some users and their correlation with other users. It therefore emphasizes that inefficiency in our model is tightly linked to data externalities.

3.7 Are Data Markets Beneficial?

Theorem 3 focuses on the comparison of the market equilibrium to the first best. This is a tough comparison for the market because in the first best some users share their data and benefit from market transactions, while others do not share. A lower bar for data markets is whether they achieve positive social surplus so that any inefficiencies they create are (partially) compensated by benefits for other agents. We next show that this is not necessarily the case and provide a

sufficient condition for the equilibrium (social) surplus to be negative — so that shutting down data markets all together would improve social surplus (and thus utilitarian welfare).

Let us also introduce the following notation: for any action profile $\mathbf{a} \in \{0, 1\}^n$, we let $\mathcal{I}_i(T)$ denote the leaked information about user i where $T = \{i \in \mathcal{V} : a_i = 1\}$.

Proposition 3. *We have*

$$\text{Social surplus}(\mathbf{a}^E) \leq \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathcal{V}) - \sum_{i \in \mathcal{V}^{(h)}} (v_i - 1) \mathcal{I}_i(\mathcal{V}^{(l)}).$$

This implies that if

$$\sum_{i \in \mathcal{V}^{(h)}} (v_i - 1) \mathcal{I}_i(\mathcal{V}^{(l)}) > \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathcal{V}), \quad (6)$$

then the equilibrium surplus is negative and utilitarian welfare improves if data markets are shut down.

This proposition follows immediately from Lemma 3. The first term is an upper bound on the gain in social surplus from the sharing decisions of low-value users (even if these gains do not necessarily accrue to the users themselves and are mainly captured by the platform). This expression is an upper bound because we are evaluating this term under the assumption that all users share their data, thus maximizing the amount of socially beneficial information about low-value users. The second term is a lower bound on the loss of privacy from high-value users. It is a lower bound because the loss of privacy is evaluated for the minimal set of agents, the low-value ones, who always share their data (in equilibrium a superset of $\mathcal{V}^{(l)}$ will share their data).

We also add that leaked information in this proposition is only a function of the matrix Σ as shown in Theorem 2, so the right-hand side is in terms of model parameters and does not depend on equilibrium objects.

The next proposition provides a sufficient condition in terms of values of privacy and correlations between data that ensures condition (6) and implies that the equilibrium necessarily has negative social surplus.

Proposition 4. *Suppose*

$$\sum_{i \in \mathcal{V}^{(h)}} \left((v_i - 1) \frac{\sum_{j \in \mathcal{V}^{(l)}} \Sigma_{ij}^2}{\|\Sigma^{(l)}\|_1 + 1} \right) > \sum_{i \in \mathcal{V}^{(l)}} \sigma_i^2 (1 - v_i), \quad (7)$$

where $\|\Sigma^{(l)}\|_1$ denotes the 1-norm of the submatrix of Σ which only includes the rows and columns corresponding to low-value users. Then the equilibrium surplus is negative.¹¹

Proposition 4 provides a sufficient condition in terms of the values of privacy and the correlation between high and low-value users for negative equilibrium surplus. It highlights the inefficiencies caused by direct data externalities which correspond to part 2 of Theorem 3. To interpret condition (7), let us fix the set of low-value users and their values. Then condition (7) is

¹¹1-norm of a matrix A is defined as $\|A\|_1 = \max_i \sum_{j=1}^n |A_{ij}|$.




Figure 3: Shaded area shows the pairs of (ρ, v_h) with negative equilibrium surplus in the setting of Example 2.

more likely to hold when there exist users with sufficiently high values and high correlation with low-value users, so that data shared by low-value users leaks a lot of information about users who value their privacy highly.

Example 2. We consider a setting with two communities, each of size 10. Suppose that all users in community 1 are low-value and have a value of privacy equal to 0.9, while all users in community 2 are high-value (with $v_h > 1$). We also take the variances of all user data to be 1, the correlation between any two users who belong to the same community to be $1/20$, and the correlation between any two users who belong to different communities to be ρ . Figure 3 depicts equilibrium surplus as a function of v_h and ρ . The curve in the figure represents the combinations of these two variables for which the social surplus is equal to zero. Moving in the northeast direction reduces equilibrium surplus and hence the shaded area has negative surplus. Consequently, in this part of the parameter space, shutting down data markets improves utilitarian social welfare. Two points are worth noting. First, relatively small values of the correlation coefficient ρ are sufficient for social surplus to be negative. Second, when v_h is very close to 1, the social surplus is always positive because the negative surplus from high-value users is compensated by the social benefits their data sharing creates for low-value users. In Example B-1 in the online Appendix, we build on this example to provide an explicit case where the platform benefits from data markets even when users lose out.

4 Competition Among Platforms

In this section we generalize the main results from the previous section to a setting in which two platforms compete for (the data of) users and focus on the case where the platforms set prices before joining decisions to attract users.¹² The timing of the events is as follows:

¹²An alternative timing of events is one where users first join platforms and then data prices are announced. In the working paper version of our work, we also analyzed this case and showed that, though the analysis is somewhat

1. Platforms simultaneously offer price vectors $\mathbf{p}^1 \in \mathbb{R}^n$ and $\mathbf{p}^2 \in \mathbb{R}^n$.
2. Users simultaneously decide which platform, if any, to join and whether to share their data.

For any $i \in \mathcal{V}$, we denote by $b_i \in \{0, 1, 2\}$ the joining decision of user i where $b_i = 0$ means user i does not join, $b_i = 1$ means she joins platform 1, and $b_i = 2$ stands for joining platform 2. Let us also define

$$J_1 = \{i \in \mathcal{V} : b_i = 1\} \text{ and } J_2 = \{i \in \mathcal{V} : b_i = 2\},$$

as the sets of users joining the two platforms.

Similar to the monopoly case in the previous section, the payoff of a platform is a function of leaked information about users and payments to users. So for platform $k \in \{1, 2\}$, we have

$$U^{(k)}(J_k, \mathbf{a}^{J_k}, \mathbf{p}^{J_k}) = \sum_{i \in J_k} \mathcal{I}_i(\mathbf{a}^{J_k}) - \sum_{i \in J_k: a_i^{J_k}=1} p_i^{J_k}, \quad (8)$$

where $\mathbf{a}^{J_k} \in \{0, 1\}^{|J_k|}$ denotes the sharing decision of users belonging to this platform, and \mathbf{p}^{J_k} denotes the vector of prices the platform offers to users in J_k .

The payoff of a user has three parts. First, each user receives a valuable service from the platform it joins. Since we are modeling joining decisions in this section, we will be more explicit about this “joining value” and assume that it depends on who else joins the platform. We therefore write this part of the payoff as $c_i(J_{b_i})$ for user i joining platform b_i , with the convention that $J_0 = \emptyset$, and also normalize $c_i(J) = 0$ for all $J \not\ni i$ and for all $i \in \mathcal{V}$. Second, the user suffers a disutility due to loss of privacy from leaked information as before, and we again denote the value of privacy for user i by v_i . Third, she receives benefits from any payments from the platform in return of the data she shares. Thus the payoff to user i joining platform $b_i \in \{1, 2\}$ is

$$u_i(a_i, b_i, \mathbf{a}_{-i}, \mathbf{b}_{-i}, \mathbf{p}^1, \mathbf{p}^2) = \begin{cases} p_i^{b_i} - v_i \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^{J_{b_i}}) + c_i(J_{b_i}) & a_i = 1 \\ -v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^{J_{b_i}}) + c_i(J_{b_i}) & a_i = 0, \end{cases} \quad (9)$$

where \mathbf{a}^{J_k} denotes the vector of sharing decisions in the set J_k for $k = 1, 2$.

Since our focus is on situations in which users join online platforms and share their data, we impose that joining values are sufficiently large.

Assumption 1. For each $i \in \mathcal{V}$, we have

1. for all J and J' such that $i \in J$ and $J \subset J'$, we have $c_i(J') > c_i(J)$.
2. $c_i(\{i\}) > \max_{j \in \mathcal{V}} v_j \sigma_j^2$.

This assumption implies that users receive greater services from a platform when there are more users on the platform, which captures the network effects in online services and social media.

simpler, similar inefficiencies apply in this case as well.

The fact that this benefit is indexed by i means that users can prefer being on the same platform with different sets of other users. The second part of this assumption imposes that the minimum value of the (non-data) services provided by the platform is larger than the maximum disutility from leaked information. In the rest of this section, we impose Assumption 1 without explicitly stating it.

Equilibria in this environment will typically be in mixed strategies, and we formally define mixed strategy equilibria in the online Appendix in terms of strategies that define probability distributions over price vectors for the platforms and user actions. Theorem B-1 there establishes the existence of a mixed strategy equilibrium with competition.¹³

The next lemma ensures that all users join one of the platforms and simplifies's are analysis in this section.¹⁴

Lemma 4. *Each user joins one of the two platforms. In other words, $b_i = 1$ or 2 for all $i \in \mathcal{V}$.*

We next show that the equilibrium is even more likely to be inefficient when platforms compete using data prices. In particular, in contrast to the settings studied so far, the equilibrium is inefficient not only when high-value users are correlated with other users, but also when there is correlation only among low-value users. For this theorem, let us define:

$$\delta = \min_{i, T \subseteq \mathcal{V}} c_i(\mathcal{V}) - c_i(T) \text{ and } \Delta = \max_{i, T \subseteq \mathcal{V}} c_i(\mathcal{V}) - c_i(T).$$

Theorem 4. 1. *Suppose every user is uncorrelated with all other users. Then the equilibrium is efficient.*

2. *Suppose that every high-value user is uncorrelated with all other users, but at least two low-value users are correlated with each other. Then there exist $\underline{\delta}$, $\bar{\Delta}$, $\tilde{\Delta}$, $\bar{\mathbf{v}}$, and $\tilde{\mathbf{v}}$ such that:*

2-1) *If $\delta \geq \underline{\delta}$, the equilibrium is efficient.*

2-2) *If $\Delta \leq \bar{\Delta}$ and $\mathbf{v}^{(l)} \leq \bar{\mathbf{v}}$, the equilibrium is efficient.*

2-3) *If $\Delta \leq \tilde{\Delta}$ and $\mathbf{v}^{(l)} \geq \tilde{\mathbf{v}}$, the equilibrium is inefficient.*

3. *Suppose that at least one high-value user is correlated with a low-value user. Then there exist $\tilde{\delta} > \bar{\Delta} > \bar{\delta} > 0$, $\bar{\mathbf{v}} \in \mathbb{R}^{|\mathcal{V}^{(h)}|}$, and $\underline{\mathbf{v}} \in \mathbb{R}^{|\mathcal{V}^{(l)}|}$ such that:*

3-1) *If $\mathbf{v}^{(h)} \geq \bar{\mathbf{v}}$, $\mathbf{v}^{(l)} \geq \underline{\mathbf{v}}$, $\Delta \leq \bar{\Delta}$, and $\delta \geq \bar{\delta}$, the equilibrium is inefficient.*

3-2) *If $\delta \geq \tilde{\delta}$, the equilibrium is efficient.*

¹³In are setting with a monopoly platform, users no longer have the option of switching to another platform, and we focus on the Stackelberg equilibrium where the platform set prices anticipating user choices and selects the most advantageous user equilibrium for itself (when there were multiple user equilibria). This ensures that an equilibrium data price yields a (weakly) greater payoff for the platform than any other price for any other user equilibrium. Because users now make their joining decisions after price offers, we use the standard Nash equilibrium notion and require that for each platform and any other price than its equilibrium price there exists a user equilibrium in which the platform's payoff is no greater than its equilibrium payoff.

¹⁴This assumption also implies that in a monopoly setting all users join the monopoly platform, which is the reason we did not introduce the joining decision in the previous section.

The first part is straightforward: without correlation there is no data externality, ensuring efficiency.

The second part is new relative to our previous results: now the equilibrium is inefficient even when high-value users are uncorrelated with all other users. This inefficiency is caused by competition using data prices. Since there is no correlation between high-value and low-value users, the first best involves all low-value users sharing their data and all (high-value and low-value) users joining the same platform in order to benefit from the highest joining values. However, we show in part 2.3 that such an allocation is not an equilibrium, because the other platform can attract some of the low-value users who can benefit by having less of their information leaked by other low-value users (even though information leakage about these users is socially beneficial, it is privately costly for them). This leads to a fragmented distribution of users across platforms, leading to inefficiency (in particular, in this case the surplus under competition is smaller than the surplus under monopoly). Parts 2.1 and 2.2 provide conditions for efficiency in terms of the c function being sufficiently steep or the privacy concerns of low-value users being sufficiently weak.

Competition affects not only efficiency but also the distribution of surplus. In particular, in the monopoly model, data prices are depressed, benefiting the platform at the expense of the consumers. Competition may partially rectify this, because low-value users may segregate between the two platforms, which reduces information leakages about them and increase data prices. Nevertheless, part 2.3 shows that this does not restore efficiency, because it fails to exploit the joining (service quality) and data-sharing benefits of having low-value users on the same platform.

Part 3.1 of the theorem is similar to our other inefficiency results. In this case, in the first best all users join the same platform (because the c function is sufficiently steep), but only low-value users uncorrelated with high-value users share their data (because $v^{(h)}$ is sufficiently high). We show, however, that this allocation cannot be an equilibrium because the other platform can deviate and attract a subset of low-value users and induce them to share their data. In part 3.2 the first best is, once again, for all users to join one of platforms (in this case, the surplus under competition is higher than the surplus under monopoly). But now because the joining values are even steeper, the other platform can no longer attract a subset of these users, while the threat of all users switching to this other platform supports the first-best allocation (though there also exist inefficient equilibria in this case). Finally, we show in the working paper version that when high-value users are uncorrelated with low-value users but correlated among themselves, the equilibrium may or may not be efficient.

5 Extensions

Our framework is purposefully stylized. This raises the question of whether some of our conclusions critically depend on our simplifying assumptions. In this section, we show that all of our main insights generalize when the correlation structure across users is more general than the Gaussian distributions we have assumed; when the value of privacy of different users are not

known by the platform(s); and when the correlation of information across users is unknown. For simplicity, we focus on the case of a monopoly platform in this section.

5.1 General Correlation Structures

As noted above, all of our results so far depend on and follow from Lemma 1, which was established using the fact that the measure of leaked information is mean square error and all data and signals are Gaussian. We also prove that this lemma holds, and all of our results readily extend, under more general assumptions so long as the following four properties hold (see Appendix A):

1. **(No leakage with independence)** If a user i 's information is independent from the information of all other users, then we have $\mathcal{I}_j(a_i = 1, \mathbf{a}_{-i}) = \mathcal{I}_j(a_i = 0, \mathbf{a}_{-i})$ for all $j \neq i$, $\mathbf{a}_{-i} \in \{0, 1\}^{n-1}$.
2. **(Leakage with non-independence)** If the information of two users i and j are non-independent (given any set of other users who share), then for any action profile $\mathbf{a} \in \{0, 1\}^n$ where user i shares her data, leaked information about user j will be non-zero. That is, $\mathcal{I}_j(a_i = 1, \mathbf{a}_{-i}) > 0$ for all $\mathbf{a}_{-i} \in \{0, 1\}^{n-1}$.
3. **(Monotonicity)** For two action profiles \mathbf{a} and \mathbf{a}' with $\mathbf{a} \geq \mathbf{a}'$, we have $\mathcal{I}_i(\mathbf{a}) \geq \mathcal{I}_i(\mathbf{a}')$ for all $i = 1, \dots, n$.
4. **(Submodularity)** For two action profiles \mathbf{a} and \mathbf{a}' with $\mathbf{a}'_{-i} \geq \mathbf{a}_{-i}$, we have $\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) \geq \mathcal{I}_i(a_i = 1, \mathbf{a}'_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}'_{-i})$.

These four properties hold under our baseline leaked information measure and Gaussian signals. We also prove in Appendix A that they and thus Lemma 1 generalize to another benchmark case: when the measure of leaked information is mutual information between a user's type and the vector of types of users who have shared their data: $\mathcal{I}_i(\mathbf{a}) = I(X_i; (X_j : a_j = 1))$ (where the mutual information between two random variables X and Y is defined as $I(X; Y) = \mathbb{E}_{X,Y}[-\log \frac{P(X)P(Y)}{P(X,Y)}]$). In particular, this measure of leaked information satisfies Properties 1-4 for any distribution of random variables \mathbf{X} .

5.2 Unknown Valuations

Our analysis has so far assumed that platforms know the value of privacy of different users. In this section, we adopt the more realistic assumption that they do not know the exact valuations of users, but understand that the value of privacy of user i , v_i , has a distribution represented by the cumulative distribution function F_i and density function f_i (with upper support denoted by v^{\max}). Users know their own value of privacy. We then show how the platform can design a mechanism to elicit this information (in the form of users reporting their value of privacy) and prove that all of the main insights from our analysis generalize to this case.

Using the revelation principle we can define an equilibrium as a pair $(\mathbf{a}^E, \mathbf{p}^E)$ of functions of the reported valuations $\mathbf{v} = (v_1, \dots, v_n)$ such that each user finds it incentive compatible to report her true value and the expected payoff of the platform is maximized taking this reporting behavior is given. That is,

$$(\mathbf{a}^E, \mathbf{p}^E) = \max_{\mathbf{a}: \mathbb{R}^n \rightarrow \{0,1\}^n, \mathbf{p}: \mathbb{R}^n \rightarrow \mathbb{R}^n} \mathbb{E}_{\mathbf{v}} \left[\sum_{i=1}^n \mathcal{I}_i(\mathbf{a}(\mathbf{v})) - \sum_{i: a_i(\mathbf{v})=1} p_i(\mathbf{v}) \right]$$

$$p_i(\mathbf{v}) - v_i \mathcal{I}_i(\mathbf{a}(\mathbf{v})) \geq p_i(\mathbf{v}_{-i}, v'_i) - v_i \mathcal{I}_i(\mathbf{a}(\mathbf{v}_{-i}, v'_i)), \quad \text{for all } v'_i, \mathbf{v}, \text{ and } i \in \mathcal{V}.$$

In the online Appendix, using an argument similar to Myerson [1981], we characterize the equilibrium in this case and prove that our inefficiency results hold with the only difference being that instead of valuations, the virtual valuations define low and high value users, where the virtual valuation of a user with value v is $\Phi_i(v) = v + \frac{F_i(v)}{f_i(v)}$.

5.3 Unknown Correlation

Another simplifying assumption we have utilized is that both the platform and the users know the correlation structure Σ . We now show that our main results generalize when this correlation structure is unknown. Suppose, in particular, that the correlation structure Σ is drawn from a distribution μ over a finite set \mathcal{S} of covariance matrices. The timing of the events is as follows. First, the platform offers prices (knowing only the distribution of correlations μ), then users decide whether they want to share their data (again knowing only the distribution of correlations μ), and finally the correlation structure is realized which together with the action profile of users determine the utility of the users and the platform. This implies that, given action profile $\mathbf{a} \in \{0, 1\}^n$, the utility of user i in this setting becomes

$$u_i(a_i, \mathbf{a}_{-i}, \mathbf{p}) = \begin{cases} p_i - v_i \mathbb{E}_{\Sigma \sim \mu} [\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i})] & a_i = 1 \\ -v_i \mathbb{E}_{\Sigma \sim \mu} [\mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})] & a_i = 0 \end{cases}$$

and the utility of the platform becomes

$$U(\mathbf{a}, \mathbf{p}) = \sum_{i \in \mathcal{V}} \mathbb{E}_{\Sigma \sim \mu} [\mathcal{I}_i(\mathbf{a})] - \sum_{i \in \mathcal{V}: a_i=1} p_i.$$

The next theorem, which is the analogue of Theorem 3, characterizes the conditions for efficiency and inefficiency in this case.¹⁵

¹⁵One may wish to go even further and investigate whether the platform can learn the distribution of valuations from past observations. This is a challenging question both because it would require an extension of our model to a fully dynamic setting and new statistical tools for learning the general variance-covariance structure from past observations.

- Theorem 5.** 1. Suppose every high-value user is uncorrelated with all other users almost surely, i.e., $\mathbb{P}_{\Sigma \sim \mu} (\Sigma_{ij} = 0) = 1$ for all $i \in \mathcal{V}^{(h)}, j \in \mathcal{V}^{(l)}$. Then the equilibrium is efficient.
2. Suppose at least one high-value user is correlated (has a non-zero correlation coefficient) with a low-value user with non-zero probability, i.e., there exists $i \in \mathcal{V}^{(h)}$ and $j \in \mathcal{V}^{(l)}$ such that $\mathbb{P}_{\Sigma \sim \mu} (\Sigma_{ij} \neq 0) > 0$. Then there exists $\bar{\mathbf{v}} \in \mathbb{R}^{|\mathcal{V}^{(h)}|}$ such that for $\mathbf{v}^{(h)} \geq \bar{\mathbf{v}}$ the equilibrium is inefficient.
3. Suppose every high-value user is uncorrelated with all low-value users almost surely and at least one high-value user is correlated with another high-value user with positive probability. Let $\tilde{\mathcal{V}}^{(h)} \subseteq \mathcal{V}^{(h)}$ be the subset of high-value users correlated with at least one other high-value user with positive probability. Then for each $i \in \tilde{\mathcal{V}}^{(h)}$ there exists $\bar{v}_i > 1$ such that if for any $i \in \tilde{\mathcal{V}}^{(h)}$ $v_i < \bar{v}_i$, the equilibrium is inefficient

6 Regulation

The inefficiencies documented so far raise the question of whether certain types of government policies or regulations could help data markets function better. We briefly address this question in this section. We first discuss taxes and then turn to a regulation scheme based on “de-correlation” to reduce the informativeness of the data of users about others. For simplicity, we focus on the case of a single platform with complete information.

6.1 Taxation

It is straightforward to establish that a simple Pigovian tax scheme, using personalized taxes on data transactions, can restore the first best (see our working paper version). This is because not taxing users who should be sharing in the first best is sufficient to ensure that they share in the post-tax equilibrium as well regardless of the sharing decisions of the rest of the users. Then imposing prohibitive taxes on the data transactions of users who should not be sharing implements the first best. Pigovian taxes implement the first best, but these taxes vary across individuals, which presupposes a huge amount of information on the part of the planner/tax authority. A natural question is whether a uniform tax scheme can also improve over the equilibrium allocation. If equilibrium surplus is negative, then a uniform and sufficiently high tax on data transactions can shut down the data market and improves equilibrium surplus. However, beyond this simple case with negative equilibrium surplus, there is no guarantee that uniform taxes on data transactions improve welfare. This is because such taxes may prevent beneficial data trades as well. We next consider an alternative regulation that keeps the beneficial data trades while eliminating the negative effects of data sharing.

6.2 Mediated Data Sharing and De-correlation

In this subsection, we propose a different approach to improving the efficiency of data markets. Our analysis has clarified that a main source of inefficiency in such markets is the correlation between the data of a user who is not sharing with the data of others who have shared their data. Our present approach is founded on the observation that the data of different users can be transformed in such a manner as to remove the correlation between any user who does not wish to share her data and all other users, while maintaining the correlation of information within the set of users sharing their data. We refer to such a scheme as *de-correlation*.

Suppose that instead of sharing their data with the platform, users share their data with a (trusted third-party) mediator, who can either not share these data with the platform (as instructed) or transform them before revealing them to the platform.¹⁶ Recall that user i 's data are represented by $S_i = X_i + Z_i$. The main idea is that the mediator collects all the data from the users and then computes transformed variables for each user removing the correlation with the information of other users and only shares the transformed data of those who are willing to sell their data (but utilizes the data of others for removing the correlation with their information).¹⁷

Formally, we consider the following de-correlation scheme: $\tilde{\mathbf{S}} = \Sigma^{-1}\mathbf{S}$ where $\mathbf{S} = (S_1, \dots, S_n)$ is the vector of data of all users. Clearly, $\tilde{\mathbf{S}}$ is jointly normal and has the property that if user i does not share her data, then the data of other users leak no information about user i 's type. This is formally stated in the next lemma.

Lemma 5. *With de-correlation, for any action profile $\mathbf{a} \in \{0, 1\}^n$, the leaked information about user i is*

$$\tilde{\mathcal{I}}_i(\mathbf{a}) = \sigma_i^2 - \min_{\hat{x}_i} \mathbb{E} \left[\left(X_i - \hat{x}_i(\tilde{\mathbf{S}}_{\mathbf{a}}) \right)^2 \right] = \begin{cases} 0 & a_i = 0 \\ \mathcal{I}_i(a_i, \mathbf{a}_{-i}) & a_i = 1. \end{cases}$$

This lemma clarifies our claim in the Introduction and shows that the de-correlation scheme leaves information leaked about the user sharing her data the same, but removes the leakage about users who are not sharing their data.

We next characterize the equilibrium pricing, denoted by $\tilde{\mathbf{p}}^E$, and sharing profile, denoted by $\tilde{\mathbf{a}}^E$, with this transformation, and show that, with de-correlation, there is no information leakage about those who do not share, and therefore they do not contribute to the platform's payoff. Moreover, the price offered to users who share must make them indifferent between sharing and not sharing and thus give them zero payoff (which they can guarantee by not sharing). Given this characterization, it follows that de-correlation always improves equilibrium surplus and, moreover, eliminates cases where the social surplus is negative.

¹⁶Obviously, a de-correlation scheme can only work if the mediator is fully reliable and trusted, and this is an important constraint in practice, which we are not dealing with in this paper.

¹⁷In practice, it may be more relevant to remove the correlation between a user's data and the average data of different user types. In that case, we can partition the set of users into K cells and apply this de-correlation procedure to the average data of cells.

Theorem 6. 1. The equilibrium sharing profile after de-correlation is given by

$$\tilde{\mathbf{a}}^E = \operatorname{argmax}_{\mathbf{a} \in \{0,1\}^n} \sum_{i \in \mathcal{V}} (1 - v_i) \tilde{\mathcal{I}}_i(\mathbf{a}),$$

with prices $\tilde{p}_i^E = v_i \tilde{\mathcal{I}}_i(\tilde{\mathbf{a}}^E)$ for any $i \in \mathcal{V}$ such that $\tilde{a}_i^E = 1$.

2. Let $(\tilde{\mathbf{a}}^E, \tilde{\mathbf{p}}^E)$ and $(\mathbf{a}^E, \mathbf{p}^E)$ denote the equilibrium with and without the de-correlation scheme, respectively. Then

$$\text{Social surplus}(\tilde{\mathbf{a}}^E) \geq \max \{ \text{Social surplus}(\mathbf{a}^E), 0 \}.$$

That equilibrium surplus increases after de-correlation is a consequence of the fact that in the original equilibrium the contribution of high-value users (who do not share) to social surplus is less than or equal to zero, while after de-correlation their contribution to social surplus is greater than or equal to zero.¹⁸ Moreover, because there are no users with negative contribution to social surplus after de-correlation, equilibrium surplus is always positive. This observation also implies that the de-correlation scheme outperforms policies that shut down data markets — since instead of achieving zero equilibrium surplus by shutting down these markets, e.g., as in Proposition 3, this scheme always guarantees positive social surplus.

Our proposed de-correlation scheme provides a simple benchmark that shows how the correlation between any user who does not wish to share her data and all other users can be removed while maintaining the socially beneficial correlation among users interested in sharing their data. A more practical version of this scheme would remove the correlation between classes of users, but still ensure that leaked information about users not wishing to share their data is minimized. An alternative regulation that may achieve similar objectives is to allow users to decide whether others' data can be used in advertisement or for the services that they receive, and this may be sufficient to remove some or all of the negative externalities. Open questions include whether de-correlation schemes and/or regulations that give additional control to users can be easily implemented and to what extent users would trust mediators or promises that others' data will not be used for obtaining information about them.

7 Conclusion

Because data generated by economic agents are useful for solving economic, social, or technical problems facing others in society and for designing or inventing new products and services, much of economic analysis in this area argues that the market may produce too little data. This paper develops the perspective that, in the presence of privacy concerns of some agents, the market may generate too much data. Moreover, because the data of a subset of users reveal information about

¹⁸As with personalized taxes, de-correlation involves a considerable amount of information being pooled in the hands of a centralized body. The difference, however, is that de-correlation, by ensuring that no information is leaked about users who do not want to share their data, makes such information pooling incentive compatible. Providing information to regulatory authorities is typically not incentive compatible.

other users, the market price of data tends to be depressed, creating the impression that users do not value their privacy much. The depressed market price of data and excessive data generation are intimately linked.

We exposit these ideas in a simple model in which a platform wishes to estimate the types of a collection of users, and each user has personal data (based on their preferences, past behavior, and contacts) which are correlated both with their type and with the data and types of other users. As a result, when a user decides to share her data with the platform, this enables the platform to improve its estimate of other users' types. We model the market for data by allowing the platform to offer prices (or other services) in exchange of data.

We prove the existence of an equilibrium in the data market and show that there will be too much data shared on the platform and the price of data will be excessively depressed. The result that the platform acquires too much data is a direct consequence of the externalities from the data of others. The root cause of depressed data prices is the submodularity of leaked information: when data sharing by other users already compromises the information of an individual, she has less incentive to protect her data and privacy. We further show that under some simple conditions the social surplus generated by data markets is negative, meaning that shutting down data markets improves (utilitarian) social welfare.

We extend these results to a setting with multiple platforms. Various different types of competition between platforms do not alter the fundamental forces leading to too much data sharing and excessively low prices of data. In fact, competition may make inefficiencies worse. This is in part because more data may be shared in the presence of competition, and also because the desire of some users to avoid excessive data sharing about them may lead to an inefficiently fragmented distribution of users across platforms, even when network externalities would be better exploited by having all users join the same platform. We also extend these results to a setting in which the value of privacy of different users are their private information.

Excessive data sharing may call for policy interventions to correct for the externalities and the excessively low prices of data. Individual-specific (Pigovian) taxes on data transactions can restore the first best. More interestingly, we propose a scheme based on mediated-data sharing that can improve welfare. In particular, in our baseline model, when equilibrium surplus is negative, shutting down data markets, for example with high uniform taxes on all data transactions, would improve welfare. But this prevents the sharing of the data of users with low value of privacy or high benefits from goods and services that depend on the platform accessing their data. We show that if user data are first shared with a mediator which transforms them before revealing them to the platform, the correlation of the data with the information of privacy-conscious users can be eliminated, and this would improve welfare relative to the option of shutting off data markets altogether.

We view our work as part of an emerging literature on data markets and the economics of privacy. Several interesting areas of research are suggested by our results. First, it is important to develop models of the marketplace for data that allow for richer types of competition between

different platforms. Second, our modeling of privacy and the use of data by the platform has been reduced-form. Distinguishing the uses of personal data for price discrimination, advertising, and designing of new products and services could lead to additional novel insights. For example, it may enable an investigation of whether applications of personal data for designing personalized services can be unbundled from their use for intrusive marketing, price discrimination, or misleading advertising. Third, there is much more to do on the effects of competition for data. One interesting direction is to allow platforms to differ in terms of the technology they use for processing data and protecting privacy, which may change the nature of competition. Finally, we only touched upon the possibility of designing new mechanisms for improving the functioning of data markets while reducing data externalities. Our proposed mechanism can be simplified and made more practical, for example, by aiming to remove the correlation between different user classes, as noted above, or by focusing on only some types of data. Other mediated data sharing arrangements or completely new approaches to this problem could be developed as well, but should take into account the possibility that third parties may not be fully trustworthy either. Finally, our result that market prices, or current user actions for protecting privacy, do not reveal the value of privacy highlights the need for careful empirical analysis documenting and estimating the value of data to platforms and the value that users attach to their privacy in the presence of data externalities.

Appendix A

In this part of the Appendix, we provide some of the proofs omitted from the text. Remaining proofs are presented in the online Appendix and the details of several of the examples discussed in the text are included in the working paper version.

Proof of Proposition 1

Recall that \mathbf{a}^W denotes the first best. For any $i \in \mathcal{V}$ we have $a_i^W = 1$ if and only if $\text{Social surplus}(\mathbf{a}_{-i}^W, a_i = 1) \geq \text{Social surplus}(\mathbf{a}_{-i}^W, a_i = 0)$. Substituting the expression for the social surplus into this equation yields

$$\sum_{j \in \mathcal{V}} (1 - v_j) (\mathcal{I}_j(\mathbf{a}_{-i}^W, a_i = 1) - \mathcal{I}_j(\mathbf{a}_{-i}^W, a_i = 0)) \geq 0. \quad (\text{A-1})$$

Conditional on the data provided by other users, i.e., $k \neq i$ for which $a_k^W = 1$, (X_j, S_i) are jointly normal and their covariance matrix is given by

$$\begin{pmatrix} \sigma_j^2 - \mathcal{I}_j(\mathbf{a}_{-i}^W, a_i = 0) & \text{Cov}(X_i, X_j \mid \mathbf{a}_{-i}^W, a_i = 0) \\ \text{Cov}(X_i, X_j \mid \mathbf{a}_{-i}^W, a_i = 0) & 1 + \sigma_i^2 - \mathcal{I}_i(\mathbf{a}_{-i}^W, a_i = 0) \end{pmatrix}.$$

Therefore, if in addition to users $k \neq i$ for which $a_k^W = 1$, user i also shares her data, then the leaked information of user j becomes

$$\mathcal{I}_j(\mathbf{a}_{-i}^W, a_i = 1) = \mathcal{I}_j(\mathbf{a}_{-i}^W, a_i = 0) + \frac{\text{Cov}(X_i, X_j \mid \mathbf{a}_{-i}^W, a_i = 0)^2}{1 + \sigma_i^2 - \mathcal{I}_i(\mathbf{a}_{-i}^W, a_i = 0)}. \quad (\text{A-2})$$

Substituting equation (A-2) into equation (A-1) completes the proof. ■

Proof of Lemma 1

Part 1, Monotonicity: In order to show that leaked information is monotonically increasing in the set of users who share, it suffices to establish that for any $i, j \in \mathcal{V}$ and $\mathbf{a}_{-j} \in \{0, 1\}^{n-1}$ we have $\mathcal{I}_i(a_j = 1, \mathbf{a}_{-j}) \geq \mathcal{I}_i(a_j = 0, \mathbf{a}_{-j})$. We next consider the two possible cases where $i = j$ and $i \neq j$ and show this inequality.

- $i = j$: conditional on shared data, the joint distribution of (X_i, S_i) is normal with covariance matrix $\begin{pmatrix} \hat{\sigma}_i^2 & \hat{\sigma}_i^2 \\ \hat{\sigma}_i^2 & 1 + \hat{\sigma}_i^2 \end{pmatrix}$, where $\hat{\sigma}_i^2 = \mathbb{E}[X_i^2 \mid \mathbf{a}_{-j}]$. We have $\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) = \sigma_i^2 - \left(\hat{\sigma}_i^2 - \frac{\hat{\sigma}_i^4}{1 + \hat{\sigma}_i^2}\right) \geq \sigma_i^2 - \hat{\sigma}_i^2 = \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$, completing the proof of this part.
- $i \neq j$: conditional on shared data, the joint distribution of (X_i, S_j) is normal with covariance matrix $\begin{pmatrix} \hat{\sigma}_i^2 & \hat{\Sigma}_{ij} \\ \hat{\Sigma}_{ij} & 1 + \hat{\sigma}_j^2 \end{pmatrix}$, where $\hat{\sigma}_i^2 = \mathbb{E}[X_i^2 \mid \mathbf{a}_{-j}]$, $\hat{\sigma}_j^2 = \mathbb{E}[X_j^2 \mid \mathbf{a}_{-j}]$, and $\hat{\Sigma}_{ij} = \mathbb{E}[X_i X_j \mid \mathbf{a}_{-j}]$. We have $\mathcal{I}_i(a_j = 1, \mathbf{a}_{-j}) = \sigma_i^2 - \left(\hat{\sigma}_i^2 - \frac{\hat{\Sigma}_{ij}^2}{1 + \hat{\sigma}_j^2}\right) \geq \sigma_i^2 - \hat{\sigma}_i^2 = \mathcal{I}_i(a_j = 0, \mathbf{a}_{-j})$, completing the proof of the monotonicity.

Part 2, Submodularity: We first introduce some additional notation for this proof. For any pair $i, j \in \mathcal{V}$, $\mathbf{a}_{-\{i,j\}}$ is the collection of all users' actions except for user i and user j . To prove this part, it suffices to establish that for any $\mathbf{a}_{-\{i,j\}} \in \{0, 1\}^{n-2}$, we have

$$\begin{aligned} & \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 1, a_i = 0) - \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 0, a_i = 0) \\ & \geq \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 1, a_i = 1) - \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 0, a_i = 1). \end{aligned}$$

Conditional on $\mathbf{a}_{-\{i,j\}}$, (X_j, S_j, S_i) has a normal distribution with covariance matrix

$$\begin{pmatrix} \hat{\sigma}_j^2 & \hat{\sigma}_j^2 & \hat{\Sigma}_{ij} \\ \hat{\sigma}_j^2 & 1 + \hat{\sigma}_j^2 & \hat{\Sigma}_{ij} \\ \hat{\Sigma}_{ij} & \hat{\Sigma}_{ij} & 1 + \hat{\sigma}_i^2 \end{pmatrix},$$

where $\hat{\sigma}_i^2 = \mathbb{E}[X_i^2 \mid \mathbf{a}_{-\{i,j\}}]$, $\hat{\sigma}_j^2 = \mathbb{E}[X_j^2 \mid \mathbf{a}_{-\{i,j\}}]$, and $\hat{\Sigma}_{ij} = \mathbb{E}[X_i X_j \mid \mathbf{a}_{-\{i,j\}}]$. Note that in writing this matrix, we are using the fact that the correlation between X_i and S_j is the same as the correlation between S_i and S_j (this holds because $S_i = X_i + Z_i$ for some independent noise Z_i).

Based on this covariance matrix,

$$\mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 1, a_i = 0) - \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 0, a_i = 0) = \frac{\hat{\sigma}_j^4}{1 + \hat{\sigma}_j^2}. \quad (\text{A-3})$$

We also have

$$\mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 1, a_i = 1) - \mathcal{I}_j(\mathbf{a}_{-\{i,j\}}, a_j = 0, a_i = 1) = \frac{\hat{\sigma}_j^4(1 + \hat{\sigma}_i^2) + \hat{\Sigma}_{ij}^2(1 + \hat{\sigma}_j^2) - 2\hat{\Sigma}_{ij}^2\hat{\sigma}_j^2}{(1 + \hat{\sigma}_i^2)(1 + \hat{\sigma}_j^2) - \hat{\Sigma}_{ij}^2} - \frac{\hat{\Sigma}_{ij}^2}{1 + \hat{\sigma}_i^2}. \quad (\text{A-4})$$

Comparing (A-3) and (A-4), the submodularity of leaked information becomes equivalent to $\hat{\sigma}_j^4(1 + \hat{\sigma}_i^2) + \hat{\Sigma}_{ij}^2(1 + \hat{\sigma}_j^2) \leq 2\hat{\sigma}_j^2(1 + \hat{\sigma}_i^2)(1 + \hat{\sigma}_j^2)$, which follows from $\hat{\Sigma}_{ij}^2 \leq \hat{\sigma}_i^2\hat{\sigma}_j^2$. ■

Proof of Lemma 2

Using Lemma 1, we first establish that the game is supermodular. The rest of the proof follows from Tarski's fixed point theorem. Specifically, for any $i \in \mathcal{V}$, we prove that the game has increasing differences property. This follows from part 2 of Lemma 1 that establishes if $\mathbf{a}'_{-i} \geq \mathbf{a}_{-i}$ then $\mathcal{I}_i(a_i = 1, \mathbf{a}'_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}'_{-i}) \leq \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$ which yields $u_i(a_i = 1, \mathbf{a}'_{-i}) - u_i(a_i = 0, \mathbf{a}'_{-i}) \geq u_i(a_i = 1, \mathbf{a}_{-i}) - u_i(a_i = 0, \mathbf{a}_{-i})$. Now consider the mapping $F : \{0, 1\}^n \rightarrow \{0, 1\}^n$ where $F_i(\mathbf{a}) = \operatorname{argmax}_{a_i \in \{0,1\}} u_i(a_i, \mathbf{a}_{-i})$. Using supermodularity of the game, this mapping is order preserving and therefore Tarski's theorem establishes that its fixed points form a complete lattice and therefore is non-empty and has greatest and least elements. Finally, note that each fixed point of the mapping F is a user equilibrium and vice versa. Therefore, the set of fixed points of the mapping F is exactly the set of user equilibria denoted by $\mathcal{A}(\mathbf{p})$. ■

Proof of Theorem 1

We prove that the following action profile and price vector constitute an equilibrium:

$$\mathbf{a}^E = \operatorname{argmax}_{\mathbf{a} \in \{0,1\}^n} \sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) + v_i \mathcal{I}_i(\mathbf{a}_{-i}, a_i = 0),$$

and $p_i^E = v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^E) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E))$, if $a_i^E = 1$ and $p_i^E = 0$ if $a_i^E = 0$. First note that $\mathbf{a}^E \in \mathcal{A}(\mathbf{p}^E)$. This is because the payoff of user i when $a_i^E = 1$ is $p_i^E - v_i \mathcal{I}_i(\mathbf{a}^E) = -v_i \mathcal{I}_i(\mathbf{a}_{-i}^E, a_i = 0)$. If user i deviates and chooses not to share, her payoff would remain unchanged. However, when $a_i^E = 0$, her payoff is $-v_i \mathcal{I}_i(\mathbf{a}_{-i}^E, a_i = 0)$, and deviation to sharing would lead to the lower payoff of $-v_i \mathcal{I}_i(\mathbf{a}_{-i}^E, a_i = 1)$. Therefore, faced with the price vector offer of \mathbf{p}^E , the users do not have a profitable deviation from \mathbf{a}^E . We next show that for any \mathbf{p} and $\mathbf{a} \in \mathcal{A}(\mathbf{p})$, we have $U(\mathbf{a}^E, \mathbf{p}^E) \geq U(\mathbf{a}, \mathbf{p})$. Since \mathbf{a} is a user equilibrium for the price vector \mathbf{p} , i.e., $\mathbf{a} \in \mathcal{A}(\mathbf{p})$, for all i such that $a_i = 1$, we must have $p_i \geq v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}))$. This is because if $p_i < v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}))$, then user i would have a profitable deviation to not

share her data. Thus,

$$\begin{aligned} U(\mathbf{a}, \mathbf{p}) &= \sum_{i \in \mathcal{V}} \mathcal{I}_i(\mathbf{a}) - \sum_{i \in \mathcal{V}: a_i=1} p_i \leq \sum_{i \in \mathcal{V}} \mathcal{I}_i(\mathbf{a}) - \sum_{i \in \mathcal{V}: a_i=1} v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})) \\ &= \sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) + v_i \mathcal{I}_i(\mathbf{a}_{-i}, a_i = 0) \leq U(\mathbf{a}^E, \mathbf{p}^E). \blacksquare \end{aligned}$$

Proof of Theorem 2

We use the following lemmas in this proof.

Lemma A-1 (Horn and Johnson [1987], Chapter 0.7). • *The inverse of a matrix in terms of its blocks*

$$\text{is } \begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} (A - BD^{-1}C)^{-1} & -A^{-1}B(D - CA^{-1}B)^{-1} \\ -D^{-1}C(A - BD^{-1}C)^{-1} & (D - CA^{-1}B)^{-1} \end{pmatrix}.$$

- *Sherman-Morrison-Woodbury formula for the inverse of rank one perturbation of matrix: Suppose $A \in \mathbb{R}^{n \times n}$ is an invertible square matrix and $u, v \in \mathbb{R}^n$ are column vectors. Then $A + uv^T$ is invertible if and only if $1 + v^T A^{-1}u \neq 0$. If $A + uv^T$ is invertible, then its inverse is $(A + uv^T)^{-1} = A^{-1} - \frac{A^{-1}uv^T A^{-1}}{1 + v^T A^{-1}u}$.*

Lemma A-2 (Feller [2008], Chapter 5, Theorem 5). *Suppose (X_1, \dots, X_n) has a normal distribution with covariance matrix Σ . The conditional distribution of X_1 given X_2, \dots, X_n is normal with covariance matrix $\Sigma_{11} - \mathbf{d}^T D^{-1} \mathbf{d}$, where D is the matrix obtained from Σ by removing the first row and the first column and $\mathbf{d} = (\Sigma_{12}, \dots, \Sigma_{1n})^T$.*

We now proceed with the proof of theorem. We first prove the existence of \mathbf{p}^a . Let $p_i^a = v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}))$. For any price vector \mathbf{p} such that $\mathbf{a} \in \mathcal{A}(\mathbf{p})$ we have

$$u_i(a_i = 1, \mathbf{a}_{-i}) = p_i - v_i \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) \geq u_i(a_i = 0, \mathbf{a}_{-i}) = -v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}), \quad \text{for all } i \text{ s.t. } a_i = 1.$$

Rearranging this inequality leads to $p_i \geq v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})) = p_i^a$. We next find the price vector \mathbf{p}^a in terms of the matrix Σ . Let $S \subseteq \{1, \dots, n\}$ be the set of users who have shared their data. Leaked information about any user i is only a function of the correlation among users in S and the correlation between user i and the users in S . The relevant covariance matrix for finding leaked information about user i is given by the rows and columns of the matrix Σ corresponding to users in $S \cup \{i\}$. Therefore, without loss of generality, we suppose that $i = 1$ and all users have shared their data and work with the entire matrix Σ . We find the equilibrium price for user 1 (the price offered to other users can be obtained similarly). With $a_1 = \dots, a_n = 1$, (X_1, S_1, \dots, S_n) is

normally distributed with covariance matrix

$$\begin{pmatrix} \sigma_1^2 & \sigma_1^2 & \Sigma_{12} & \dots & \Sigma_{1n} \\ \sigma_1^2 & 1 + \sigma_1^2 & \Sigma_{12} & \dots & \Sigma_{1n} \\ \Sigma_{12} & \Sigma_{12} & 1 + \sigma_2^2 & \dots & \Sigma_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Sigma_{1n} & \Sigma_{1n} & \Sigma_{2n} & \dots & 1 + \sigma_n^2 \end{pmatrix}.$$

Therefore, using Lemma A-2, the conditional distribution of X_1 given s_1, \dots, s_n is normal with variance $\sigma_1^2 - (\sigma_1^2, \Sigma_{12}, \dots, \Sigma_{1n})(I + \Sigma)^{-1}(\sigma_1^2, \Sigma_{12}, \dots, \Sigma_{1n})^T$. The best estimator of X_1 given s_1, \dots, s_n is its mean which leads to the following leaked information

$$\mathcal{I}_1(a_1 = 1, \mathbf{a}_{-1}) = (\sigma_1^2, \Sigma_{12}, \dots, \Sigma_{1n})(I + \Sigma)^{-1}(\sigma_1^2, \Sigma_{12}, \dots, \Sigma_{1n})^T. \quad (\text{A-5})$$

If user 1 deviates to $a_1 = 0$, then (X_1, S_2, \dots, S_n) has a normal distribution with covariance

$$\begin{pmatrix} \sigma_1^2 & \Sigma_{12} & \dots & \Sigma_{1n} \\ \Sigma_{12} & \ddots & & \vdots \\ \vdots & & I + D & \\ \Sigma_{1n} & \dots & & \ddots \end{pmatrix},$$

where D is obtained from Σ by removing the first row and column. Therefore, using Lemma A-2, the conditional distribution of X_1 given s_2, \dots, s_n is normal with variance $\sigma_1^2 - (\Sigma_{12}, \dots, \Sigma_{1n})(I + D)^{-1}(\Sigma_{12}, \dots, \Sigma_{1n})^T$ and leaked information of user 1 is

$$\mathcal{I}_1(a_1 = 0, \mathbf{a}_{-1}) = (\Sigma_{12}, \dots, \Sigma_{1n})(I + D)^{-1}(\Sigma_{12}, \dots, \Sigma_{1n})^T. \quad (\text{A-6})$$

Using A-5 and A-6, the price offered to user 1 must satisfy $\frac{p_1^a}{v_1} = (\sigma_1^2, \mathbf{d}^T)^T \begin{pmatrix} \sigma_1^2 + 1 & \mathbf{d}^T \\ \mathbf{d} & (I + D) \end{pmatrix}^{-1} (\sigma_1^2, \mathbf{d}^T) - \mathbf{d}^T (I + D)^{-1} \mathbf{d}$, where $\mathbf{d} = (\Sigma_{12}, \dots, \Sigma_{1n})$. We next simplify the right-hand side of the above equation. Using part 1 of Lemma A-1,

$$(\sigma_1^2, \mathbf{d}^T)^T \begin{pmatrix} \sigma_1^2 + 1 & \mathbf{d}^T \\ \mathbf{d} & I + D \end{pmatrix}^{-1} (\sigma_1^2, \mathbf{d}^T) - \mathbf{d}^T (I + D)^{-1} \mathbf{d} = (\sigma_1^2, \mathbf{d}^T)^T M (\sigma_1^2, \mathbf{d}^T) - \mathbf{d}^T (I + D)^{-1} \mathbf{d},$$

with

$$M = \begin{pmatrix} ((\sigma_1^2 + 1) - \mathbf{d}^T (I + D)^{-1} \mathbf{d})^{-1} & -\frac{1}{\sigma_1^2 + 1} \mathbf{d}^T \left((I + D) - \frac{1}{1 + \sigma_1^2} \mathbf{d} \mathbf{d}^T \right)^{-1} \\ -(I + D)^{-1} \mathbf{d} \left((\sigma_1^2 + 1) - \mathbf{d}^T (I + D)^{-1} \mathbf{d} \right)^{-1} & \left((I + D) - \frac{1}{1 + \sigma_1^2} \mathbf{d} \mathbf{d}^T \right)^{-1} \end{pmatrix}.$$

Using part 2 of Lemma A-1 and $\mathcal{I}_1(a_1 = 0, \mathbf{a}_{-1}) = \mathbf{d}^T (I + D)^{-1} \mathbf{d}$, we can further simplify this

equation to $\frac{(\sigma_1^2 - \mathcal{I}_1(a_1=0, \mathbf{a}_{-1}))^2}{(\sigma_1^4 + 1) - \mathcal{I}_1(a_1=0, \mathbf{a}_{-1})}$. This also implies the decomposition stated in the theorem. ■

Proof of Corollary 1

Using Theorem 2, we have $p_i^{(a_i=1, \mathbf{a}_{-i})} = v_i \frac{(\sigma_i^2 - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}))^2}{(\sigma_i^4 + 1) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i})}$, which is increasing in σ_i^2 and decreasing in $\mathcal{I}_i(a_i=0, \mathbf{a}_{-i})$. Again, using Theorem 2, we have $\mathcal{I}_i(a_i=1, \mathbf{a}_{-i}) = \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}) + \frac{(\sigma_i^2 - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}))^2}{(\sigma_i^4 + 1) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i})}$, which is increasing in both $\mathcal{I}_i(a_i=0, \mathbf{a}_{-i})$ and σ_i^2 . ■

Proof of Proposition 2

Let $i \in \mathcal{V}$ be such that $a'_i = a_i = 1$. Using Theorem 2, we have $p_i^{\mathbf{a}} = v_i (\mathcal{I}_i(a_i=1, \mathbf{a}_{-i}) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i})) \stackrel{(a)}{\geq} v_i (\mathcal{I}_i(a'_i=1, \mathbf{a}'_{-i}) - \mathcal{I}_i(a'_i=0, \mathbf{a}'_{-i})) = p_i^{\mathbf{a}'}$, where (a) follows from submodularity of leaked information, i.e., part 2 of Lemma 1. ■

Proof of Lemma 3

Suppose to obtain a contradiction that in equilibrium $a_i^E = 0$ for some $i \in \mathcal{V}$ with $v_i \leq 1$. We prove that there exists a deviation which increases the platform's payoff. In particular, the platform can deviate and offer price $p_i = v_i (\mathcal{I}_i(a_i=1, \mathbf{a}_{-i}^E) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E))$ so that user i shares.

From Theorem 1, the equilibrium action profile \mathbf{a}^E must maximize $\sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) + v_i \mathcal{I}_i(a_i=0, \mathbf{a}_{-i})$. We show that $(a_i=1, \mathbf{a}_{-i}^E)$ increases this objective, which yields a contradiction:

$$\begin{aligned} & \left(\sum_{j \in \mathcal{V} \setminus \{i\}} \mathcal{I}_j(a_j=1, \mathbf{a}_{-j}^E) - \mathcal{I}_j(a_j=0, \mathbf{a}_{-j}^E) \right) - \left(\sum_{j \in \mathcal{V}: a_j^E=1} p_j^{(a_i=1, \mathbf{a}_{-i}^E)} - p_j^{(a_i=0, \mathbf{a}_{-i}^E)} \right) \\ & + ((1 - v_i) \mathcal{I}_i(a_i=1, \mathbf{a}_{-i}^E) + v_i \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E)) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E) \\ & \stackrel{(a)}{\geq} - \left(\sum_{j \in \mathcal{V}: a_j^E=1} p_j^{(a_i=1, \mathbf{a}_{-i}^E)} - p_j^{(a_i=0, \mathbf{a}_{-i}^E)} \right) + ((1 - v_i) \mathcal{I}_i(a_i=1, \mathbf{a}_{-i}^E) + v_i \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E)) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E) \\ & \stackrel{(b)}{\geq} (1 - v_i) (\mathcal{I}_i(a_i=1, \mathbf{a}_{-i}^E) - \mathcal{I}_i(a_i=0, \mathbf{a}_{-i}^E)) \stackrel{(c)}{\geq} 0, \end{aligned}$$

where (a) follows from monotonicity of leaked information (i.e., part 1 of Lemma 1), (b) follows from Proposition 2, and (c) follows from the fact that $v_i \leq 1$ and leaked information is monotone. This shows that for any i such that $v_i \leq 1$ we must have $a_i^E = 1$. ■

Proof of Theorem 3

We use the following notation in this proof. For any action profile $\mathbf{a} \in \{0, 1\}^n$ and any subset $T \subseteq \{1, \dots, n\}$, we let \mathbf{a}_T denote a vector that include all the entries of a_i for which $i \in T$.

Part 1: For a given action profile \mathbf{a} , the social surplus can be written as

$$\begin{aligned} \text{Social surplus}(\mathbf{a}) &= \sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) \\ &\stackrel{(a)}{=} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}_{\mathcal{V}^{(l)}}, \mathbf{a}_{\mathcal{V}^{(h)}} = \mathbf{0}) + \sum_{i \in \mathcal{V}^{(h)}} (1 - v_i) \mathcal{I}_i(a_i, \mathbf{a}_{-i} = \mathbf{0}) \stackrel{(b)}{\leq} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathcal{V}^{(l)}), \end{aligned}$$

where (a) follows from the fact that the data of high-value users are not correlated with the data of any other user, and (b) follows from the fact that for $i \in \mathcal{V}^{(l)}$, leaked information about user i (weakly) increases in the set of users who share (from part 1 of Lemma 1) and $1 - v_i \geq 0$. Conversely, for $i \in \mathcal{V}^{(h)}$ we have $1 - v_i < 0$. This implies $\mathbf{a}_i^W = 1$ if and only if $i \in \mathcal{V}^{(l)}$.

The payoff of the platform for a given action profile \mathbf{a} (and the corresponding equilibrium prices to sustain it) can be written as

$$\begin{aligned} U(\mathbf{a}, \mathbf{p}^{\mathbf{a}}) &= \sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) + v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}) \\ &\stackrel{(a)}{=} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}_{\mathcal{V}^{(l)}}) + v_i \mathcal{I}_i(\mathbf{a}_{\mathcal{V}^{(l)} \setminus \{i\}}) + \sum_{i \in \mathcal{V}^{(h)}} (1 - v_i) \mathcal{I}_i(a_i, \mathbf{a}_{-i} = \mathbf{0}) \\ &\stackrel{(b)}{\leq} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}_{\mathcal{V}^{(l)}}) + v_i \mathcal{I}_i(\mathbf{a}_{\mathcal{V}^{(l)} \setminus \{i\}}) \stackrel{(c)}{\leq} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathcal{V}^{(l)}) + v_i \mathcal{I}_i(\mathcal{V}^{(l)} \setminus \{i\}) \end{aligned}$$

where (a) follows from the fact that the data of high-value users are not correlated with the data of any other user, (b) follows from the fact that $1 - v_i < 0$ for $i \in \mathcal{V}^{(h)}$, and (c) follows from Lemma 3. Therefore, no high-value user shares in equilibrium and we have $\mathbf{a}^E = \mathbf{a}^W$.

Part 2: Let $i \in \mathcal{V}^{(l)}$ and $j \in \mathcal{V}^{(h)}$ be such that $\Sigma_{ij} > 0$. Therefore, there exists $\delta > 0$ such that $\mathcal{I}_j(\mathcal{V}^{(l)}) = \delta > 0$. We next show that for $v_j > 1 + \frac{\sum_{i \in \mathcal{V}^{(l)}} \sigma_i^2}{\delta}$ the surplus of the action profile \mathbf{a}^E is negative, establishing that it does not coincide with the first best. We have

$$\begin{aligned} \text{Social surplus}(\mathbf{a}^E) &= \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}^E) + \sum_{i \in \mathcal{V}^{(h)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}^E) \stackrel{(a)}{\leq} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \sigma_i^2 + \sum_{i \in \mathcal{V}^{(h)}} (1 - v_i) \mathcal{I}_i(\mathbf{a}^E) \\ &\stackrel{(b)}{\leq} \left(\sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \sigma_i^2 \right) + (1 - v_j) \mathcal{I}_j(\mathcal{V}^{(l)}) \leq \left(\sum_{i \in \mathcal{V}^{(l)}} \sigma_i^2 \right) + (1 - v_j) \mathcal{I}_j(\mathcal{V}^{(l)}) \stackrel{(c)}{<} 0 \end{aligned}$$

where in (a) for low-value users we have upper bounded leaked information with its maximum; in (b) we removed all the negative terms in the second summation except for the one corresponding to j for which we replaced the leaked information (of equilibrium action profile) by its minimum (using Lemma 3); and in (c) we used $v_j > 1 + \frac{\sum_{i \in \mathcal{V}^{(l)}} \sigma_i^2}{\delta}$.

Part 3: Let $i, k \in \mathcal{V}^{(h)}$ be such that $\Sigma_{ik} > 0$. The first best involves all low-values are sharing their data and none of the high-value users doing so. We next show that if the value of privacy for high-value user i is small enough, then at least one high-value user shares in equilibrium. We show this

by assuming the contrary and then reaching a contradiction. Suppose that none of high-value users share. We show that if user i shares, the platform's payoff increases. We let \mathbf{a}^m denote the sharing profile in which all users in $\mathcal{V}^{(l)} \cup \{i\}$ share their data and $\mathbf{a} \in \{0, 1\}^n$ denote the sharing profile in which all users in $\mathcal{V}^{(l)}$ share their data. Using this notation, let us write

$$\begin{aligned} U(\mathbf{a}', \mathbf{p}^{\mathbf{a}'}) &= (1 - v_i)\mathcal{I}_i(\mathcal{V}^{(l)} \cup \{i\}) + v_i\mathcal{I}_i(\mathcal{V}^{(l)}) + \sum_{k \in \mathcal{V}^{(h)} \setminus \{i\}} \mathcal{I}_k(\mathcal{V}^{(l)} \cup \{i\}) \\ &+ \left(\sum_{j \in \mathcal{V}^{(l)}} (1 - v_j)\mathcal{I}_j(\mathcal{V}^{(l)} \cup \{i\}) + v_j\mathcal{I}_j(\mathcal{V}^{(l)} \cup \{i\} \setminus \{j\}) \right) \\ &\stackrel{(a)}{=} (1 - v_i)\mathcal{I}_i(\mathcal{V}^{(l)} \cup \{i\}) + v_i\mathcal{I}_i(\mathcal{V}^{(l)}) + \left(\sum_{k \in \mathcal{V}^{(h)} \setminus \{i\}} \mathcal{I}_k(\mathcal{V}^{(l)} \cup \{i\}) \right) + U(\mathbf{a}, \mathbf{p}^{\mathbf{a}}) \stackrel{(b)}{>} U(\mathbf{a}, \mathbf{p}^{\mathbf{a}}), \end{aligned}$$

where (a) follows from the fact that high- and low-value users are uncorrelated and (b) follows by letting $v_i < \frac{\mathcal{I}_i(\mathcal{V}^{(l)} \cup \{i\}) + \sum_{k \in \mathcal{V}^{(h)} \setminus \{i\}} \mathcal{I}_k(\mathcal{V}^{(l)} \cup \{i\})}{\mathcal{I}_i(\mathcal{V}^{(l)} \cup \{i\}) - \mathcal{I}_i(\mathcal{V}^{(l)})} = \frac{\mathcal{I}_i(\{i\}) + \sum_{k \in \mathcal{V}^{(h)} \setminus \{i\}} \mathcal{I}_k(\{i\})}{\mathcal{I}_i(\{i\})}$. Finally, note that using $\Sigma_{ik} > 0$, the right-hand side of the above inequality is strictly larger than 1. The proof is completed by letting $\bar{v}_i = \frac{\mathcal{I}_i(\{i\}) + \sum_{k \in \mathcal{V}^{(h)} \setminus \{i\}} \mathcal{I}_k(\{i\})}{\mathcal{I}_i(\{i\})}$. ■

Proof of Proposition 3

For an equilibrium action profile \mathbf{a}^E , social surplus is $\text{Social surplus}(\mathbf{a}^E) = \sum_{i \in \mathcal{V}} (1 - v_i)\mathcal{I}_i(\mathbf{a}^E) \stackrel{(a)}{\leq} \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i)\mathcal{I}_i(\mathcal{V}) + \sum_{i \in \mathcal{V}^{(h)}} (1 - v_i)\mathcal{I}_i(\mathcal{V}^{(l)})$ where (a) follows from the fact that for $i \in \mathcal{V}^{(l)}$, leaked information about user i increases in the set of users who share (i.e., part 1 of Lemma 1) and $1 - v_i \geq 0$; and for $i \in \mathcal{V}^{(h)}$ we have $1 - v_i < 0$ and $\mathcal{I}_i(\mathbf{a}^E) \geq \mathcal{I}_i(\mathcal{V}^{(l)})$ by using Lemma 3. ■

Proof of Proposition 4

Using Theorem 2, leaked information about a high-value user $i \in \mathcal{V}^{(h)}$ if low-value users share is

$$(\Sigma_{ij_1}, \dots, \Sigma_{ij_k})((I + \Sigma) + M)^{-1}(\Sigma_{ij_1}, \dots, \Sigma_{ij_k})^T,$$

where low-value users are denoted by j_1, \dots, j_k and the diagonal entries of M are zero and $M_{r,s}$ is the covariance between two low-value users r and s . We next prove that this leaked information is larger than or equal to $\sum_{l=1}^k \frac{\Sigma_{ij_l}^2}{\|\Sigma\|_1 + 1}$, where $\|\Sigma\|_1 = \max_{i=1, \dots, n} \sum_{j=1}^n |\Sigma_{ij}|$. We first show that $((I + \Sigma) + M)^{-1} - ((\|\Sigma\|_1 + 1)I)^{-1} \succeq 0$ (i.e., the matrix $((I + \Sigma) + M)^{-1} - ((\|\Sigma\|_1 + 1)I)^{-1}$ is positive semidefinite). Letting μ_i denote an eigen value of the matrix $((I + \Sigma) + M)^{-1} - ((\|\Sigma\|_1 + 1)I)^{-1}$, it suffices to show that $\mu_i \geq 0$. There exists an eigenvalue, λ_i , of the matrix $(I + \Sigma) + M$ for which we have $\mu_i = \frac{1}{\lambda_i} - \frac{1}{\|\Sigma\|_1 + 1}$. We next show that all eigenvalues of the matrix $(I + \Sigma) + M$ are (weakly) smaller than $\|\Sigma\|_1 + 1$, which establishes that $\mu_i \geq 0$. Using Gershgorin Circle Theorem, the matrix $(\|\Sigma\|_1 + 1)I - ((I + \Sigma) + M)$ is positive semidefinite. This is because for row i of this

matrix, the diagonal entry is $\|\Sigma\|_1 - \Sigma_{ii}$ which is larger than the summation of the absolute values of the off-diagonal entries $\sum_{j \neq i} \Sigma_{ij}$. Therefore, for any eigenvalue of the matrix $(I + \Sigma) + M$ such as λ_i , we have $\lambda_i \leq \|\Sigma\|_1 + 1$. We can write

$$\begin{aligned} & (\Sigma_{ij_1}, \dots, \Sigma_{ij_k})((I + \Sigma) + M)^{-1}(\Sigma_{ij_1}, \dots, \Sigma_{ij_k})^T \\ & \geq (\Sigma_{ij_1}, \dots, \Sigma_{ij_k})(\|\Sigma\|_1 + 1)I^{-1}(\Sigma_{ij_1}, \dots, \Sigma_{ij_k})^T = \sum_{l=1}^k \frac{\Sigma_{ij_l}^2}{\|\Sigma\|_1 + 1}. \end{aligned} \quad (\text{A-7})$$

Using Proposition 3, equilibrium surplus is negative if $\sum_{i \in \mathcal{V}^{(h)}} (v_i - 1) \mathcal{I}_i(\mathcal{V}^{(l)}) > \sum_{i \in \mathcal{V}^{(l)}} (1 - v_i) \mathcal{I}_i(\mathcal{V})$. From inequality (A-7) and $\mathcal{I}_i(\mathcal{V}) \leq \sigma_i^2$, this condition holds provided that $\sum_{i \in \mathcal{V}^{(h)}} (v_i - 1) \frac{\sum_{j \in \mathcal{V}^{(l)}} \Sigma_{ij}^2}{\|\Sigma^{(l)}\|_1 + 1} > \sum_{i \in \mathcal{V}^{(l)}} \sigma_i^2 (1 - v_i)$, where $\Sigma^{(l)}$ denotes the submatrix of Σ which only includes the rows and columns corresponding to low value users. This completes the proof. ■

Proof of Generalization of Lemma 3 under Properties 1-4

The proof follows the proof of Lemma 3 closely, and we provide a sketch, emphasizing the places where we use Properties 1-4. Suppose to obtain a contradiction that in equilibrium $a_i^E = 0$ for some $i \in \mathcal{V}$ with $v_i \leq 1$. The equilibrium action profile \mathbf{a}^E must maximize $\sum_{i \in \mathcal{V}} (1 - v_i) \mathcal{I}_i(\mathbf{a}) + v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i})$. We show that $(a_i = 1, \mathbf{a}_{-i}^E)$ increases this objective, which yields a contradiction:

$$\begin{aligned} & \left(\sum_{j \in \mathcal{V} \setminus \{i\}} \mathcal{I}_j(a_i = 1, \mathbf{a}_{-i}^E) - \mathcal{I}_j(a_i = 0, \mathbf{a}_{-i}^E) \right) - \left(\sum_{j \in \mathcal{V}: a_j^E = 1} p_j^{(a_i=1, \mathbf{a}_{-i}^E)} - p_j^{(a_i=0, \mathbf{a}_{-i}^E)} \right) \\ & + ((1 - v_i) \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^E) + v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E)) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E) \\ & \stackrel{(a)}{\geq} - \left(\sum_{j \in \mathcal{V}: a_j^E = 1} p_j^{(a_i=1, \mathbf{a}_{-i}^E)} - p_j^{(a_i=0, \mathbf{a}_{-i}^E)} \right) + ((1 - v_i) \mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^E) + v_i \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E)) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E) \\ & \stackrel{(b)}{\geq} (1 - v_i) (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^E) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E)) \stackrel{(c)}{\geq} 0, \end{aligned}$$

where (a) follows from monotonicity of leaked information (i.e., Property 3), (b) follows from the fact that the price is $p_i = v_i (\mathcal{I}_i(a_i = 1, \mathbf{a}_{-i}^E) - \mathcal{I}_i(a_i = 0, \mathbf{a}_{-i}^E))$ and the submodularity of leaked information (i.e., Property 4), and (c) follows from the fact that $v_i \leq 1$ and leaked information is monotone (i.e., Property 3). This shows that for any i such that $v_i \leq 1$ we must have $a_i^E = 1$. The proof of this lemma uses properties 3 and 4. A generalization of Theorem 3 under properties 1-4 uses this lemma and properties 1 and 2. ■

Mutual Information Satisfies Properties 1-4

No Leakage with independence: If X_i is independent from $(X_j : j \neq i)$, then for any action profile \mathbf{a}_{-i} and any user j we have $\mathcal{I}_j(a_i = 1, \mathbf{a}_{-i}) = I(X_j; (X_i, Y)) \stackrel{(a)}{=} H(Y) + H(X_i|Y) -$

$(H(Y|X_j) + H(X_i|X_j, Y)) \stackrel{(b)}{=} H(Y) + H(X_i) - H(Y|X_j) - H(X_i) = H(Y) - H(Y|X_j) = I(X_j; Y) = \mathcal{I}_j(a_i = 0, \mathbf{a}_{-i})$, where $Y = (X_k : a_k = 1, k \neq i, j)$; (a) follows from the definition of mutual information and the entropy function and the chain rule for entropy; and (b) follows from the fact that X_i is independent of the rest of the random variables.

Leakage with non-independence: If X_i and X_j are non-independent conditional on any other set of random variables, then for any action profile \mathbf{a}_{-i} we have $\mathcal{I}_j(a_i = 1, \mathbf{a}_{-i}) - \mathcal{I}_j(a_i = 0, \mathbf{a}_{-i}) = I(X_j; (X_i, Y)) - I(X_j; Y) = I(X_i; X_j|Y) \stackrel{(a)}{>} 0$, where $Y = (X_k : a_k = 1, k \neq i, j)$, and (a) follows from the fact that mutual information is non-negative and becomes zero if and only if the two random variables are independent.

Monotonicity: For $i \in \mathcal{V}$, let $Y = (X_j : a_j = 1)$ and $Z = (X_j : a'_j = 1, a_j = 0)$. Then the inequality $I_i(\mathbf{a}') \geq I_i(\mathbf{a})$ becomes equivalent to $I(X_i; Y, Z) \geq I(X_i; Y)$. This inequality holds because $I(X_i; Y, Z) = I(X_i; Y) + I(X_i; Z|Y) \geq I(X_i; Y)$, where we used the chain rule for mutual information and positivity of (conditional) mutual information (see e.g., [Cover and Thomas \[2012\]](#)).

Submodularity: For $i \in \mathcal{V}$, let $Y = (X_j : a_j = 1, j \neq i)$ and $Z = (X_j : a'_j = 1, a_j = 0, j \neq i)$. Then the inequality $I_i(a_i = 1, \mathbf{a}_{-i}) - I_i(a_i = 0, \mathbf{a}_{-i}) \geq I_i(a_i = 1, \mathbf{a}'_{-i}) - I_i(a_i = 0, \mathbf{a}'_{-i})$ becomes equivalent to $I(X_i; X_i, Y) - I(X_i; Y) \geq I(X_i; X_i, Y, Z) - I(X_i; Y, Z)$, that in turn is equivalent to $I(X_i; Y, Z) \geq I(X_i; Y)$, which we already established above.

References

- A. Acquisti and H. R. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3): 367–381, 2005.
- A. Acquisti, C. Taylor, and L. Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–92, 2016.
- A. R. Admati and P. Pfleiderer. A monopolistic market for information. *Journal of Economic Theory*, 39(2):400–438, 1986.
- A. Agarwal, M. Dahleh, and T. Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 701–726, 2019.
- A. Agrawal, J. Gans, and A. Goldfarb. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press, 2018.
- J. J. Anton and D. A. Yao. The sale of ideas: Strategic disclosure, property rights, and contracting. *Review of Economic Studies*, 69(3):513–531, 2002.
- S. Athey, C. Catalini, and C. Tucker. The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research, 2017.
- M. Babaioff, R. Kleinberg, and R. Paes Leme. Optimal mechanisms for selling information. In *Proceedings of the 2012 ACM Conference on Electronic Commerce*, pages 92–109, 2012.

- J. Begenau, M. Farboodi, and L. Veldkamp. Big data in finance and the growth of large firms. *Journal of Monetary Economics*, 97:71–87, 2018.
- D. Bergemann and A. Bonatti. Selling cookies. *American Economic Journal: Microeconomics*, 7(3): 259–94, 2015.
- D. Bergemann and A. Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11, 2019.
- D. Bergemann, A. Bonatti, and A. Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018.
- D. Bergemann, A. Bonatti, and T. Gan-a. Markets for information. *mimeo, MIT and Yale*, 2019.
- D. Boyd. Networked privacy. *Personal Democracy Forum*. New York, NY, 2011.
- M. Burkschat and N. Torrado. On the reversed hazard rate of sequential order statistics. *Statistics & Probability Letters*, 85:106–113, 2014.
- Y. Chen and S. Zheng. Prior-free data acquisition for accurate statistical estimation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 659–677, 2019.
- Y. Chen, N. Immorlica, B. Lucier, V. Syrgkanis, and J. Ziani. Optimal data acquisition for statistical estimation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 27–44, 2018.
- J. P. Choi, D.-S. Jeon, and B.-C. Kim. Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173:113–124, 2019.
- E. H. Clarke. Multipart pricing of public goods. *Public choice*, 11(1):17–33, 1971.
- T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- J. Crémer and R. P. Mclean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.
- P. Dasgupta, E. Maskin, et al. The existence of equilibrium in discontinuous economic games, i: Theory. *Review of Economic Studies*, 53(1):1–26, 1986.
- C. Dwork, A. Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- K. Eliaz, R. Eilat, and X. Mu. Optimal privacy-constrained mechanisms. *Available at SSRN: <https://ssrn.com/abstract=3336795>*, 2019.
- P. Esó and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.
- I. P. Fainmesser, A. Galeotti, and R. Momot. Digital privacy. *Available at SSRN: <https://ssrn.com/abstract=3459274>*, 2019.
- J. A. Fairfield and C. Engel. Privacy as a public good. *Duke LJ*, 65:385, 2015.
- M. Farboodi, R. Mihet, T. Philippon, and L. Veldkamp. Big data and firm dynamics. In *AEA papers and proceedings*, volume 109, pages 38–42, 2019.

- W. Feller. *An introduction to probability theory and its applications*, volume 2. John Wiley & Sons, 2008.
- A. Ghosh and A. Roth. Selling privacy at auction. *Games and Economic Behavior*, 91:334–346, 2015.
- A. Goldfarb and C. Tucker. Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404, 2011.
- A. Goldfarb and C. Tucker. Shifts in privacy concerns. *American Economic Review*, 102(3):349–53, 2012.
- R. Gradwohl. Information sharing and privacy in networks. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 349–350, 2017.
- T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
- R. A. Horn and C. R. Johnson. Matrix analysis. Cambridge, 1987.
- J. Horner and A. Skrzypacz. Selling information. *Journal of Political Economy*, 124(6):1515–1562, 2016.
- S. Ichihashi. Non-competing data intermediaries. Available at SSRN: <https://ssrn.com/abstract=3310410>, 2019.
- S. Ichihashi. Dynamic privacy choices. In *Proceedings of the 2020 ACM Conference on Economics and Computation*, pages 539–540, 2020a.
- S. Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95, 2020b.
- C. I. Jones, C. Tonetti, et al. Nonrivalry and the economics of data. In *2018 Meeting Papers*, number 477. Society for Economic Dynamics, 2018.
- B. Jullien, Y. Lefouili, and M. H. Riordan. Privacy protection, security, and consumer retention. *Security, and Consumer Retention (June 1, 2020)*, 2020.
- K. C. Laudon. Markets and privacy. *Communications of the ACM*, 39(9):92–104, 1996.
- M. MacCarthy. New directions in privacy: Disclosure, unfairness and externalities. Available at SSRN: <https://ssrn.com/abstract=3093301>.
- R. Montes, W. Sand-Zantman, and T. Valletti. The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3):1342–1362, 2019.
- R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- F. Pasquale. *The black box society*. Harvard University Press, 2015.
- E. A. Posner and E. G. Weyl. *Radical markets: Uprooting capitalism and democracy for a just society*. Princeton University Press, 2018.
- R. A. Posner. The economics of privacy. *American Economic Review*, 71(2):405–409, 1981.
- G. J. Stigler. An introduction to privacy in economics and politics. *The Journal of Legal Studies*, pages 623–644, 1980.

- C. R. Taylor. Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650, 2004.
- J. Tirole. Digital dystopia. 2019.
- H. R. Varian. Economic aspects of personal privacy. In *Internet policy and economics*, pages 101–109. Springer, 2009.
- L. Veldkamp and C. Chung. Data and the aggregate economy. *Annual Meeting Plenary (No. 2019-1)*, *Society for Economic Dynamics*, 2019.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- S. D. Warren and L. D. Brandeis. The right to privacy. *Harvard law review*, pages 193–220, 1890.
- A. F. Westin. Privacy and freedom. *Washington and Lee Law Review*, 25(1):166, 1968.
- S. Zuboff. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books, 2019.